



ELSEVIER

Contents lists available at ScienceDirect

Journal of International Money and Finance

journal homepage: www.elsevier.com/locate/jimf



Heterogeneous agents, the financial crisis and exchange rate predictability [☆]



Daniel Buncic ^{a,*}, Gion Donat Piras ^b

^a *Institute of Mathematics & Statistics, University of St. Gallen, St. Gallen, Switzerland*

^b *KPMG, Zürich, Switzerland*

ARTICLE INFO

Article history:

Available online 25 September 2015

JEL Classification:

C22

C52

C53

E17

F31

G17

Keywords:

Empirical heterogeneous agent model

Forecasting

State-space modelling

Model combination

Exchange rate predictability

Financial crisis

ABSTRACT

We construct an empirical heterogeneous agent model which optimally combines forecasts from fundamentalist and chartist agents and evaluates its out-of-sample forecast performance using daily data covering an overall period from January 1999 to June 2014 for six of the most widely traded currencies. We use daily financial data such as level, slope and curvature yield curve factors, equity prices, as well as risk aversion and global trade activity measures in the fundamentalist agent's predictor set to obtain a proxy for the market's view on the state of the macroeconomy. Chartist agents rely upon standard momentum, moving average and relative strength index technical indicators in their predictor set. Individual agent specific forecasts are constructed using a flexible dynamic model averaging framework and are then aggregated into a model combined forecast using a forecast combination regression. We show that our empirical heterogeneous agent model produces statistically significant and economically sizeable forecast improvements over a random walk benchmark.

© 2015 Elsevier Ltd. All rights reserved.

[☆] We are grateful to Andrea Vedolin (LSE) for providing the yield curve factor loadings for Japan and Switzerland. We would like to thank Adrian Pagan, Paul Söderlind, Michael Bauer, Valentyn Panchenko, Andreas Schrimpf, George Tauchen, Graham Elliott, David Rapach, Lorenzo Camponovo, Alexander Apokin, Roland Meeks, Philip Vermeulen, Mikhail Oet, and Huichou Huang, as well as seminar and conference participants at the University of St. Gallen, University of Graz, the International Symposium on Econometric Theory and Applications in Tokyo (2015), the International Symposium on Forecasting in Riverside (2015), the annual conference of the International Association for Applied Econometrics in Thessaloniki (2015), and the annual meeting of the Financial Management Association in Orlando (2015) for helpful comments and discussions.

* Corresponding author. Tel.: +41 (71) 224 2604.

E-mail address: daniel.buncic@unisg.ch (D. Buncic).

1. Introduction

Since the seminal work of Meese and Rogoff (1983), it is well known that standard macroeconomic models of exchange rate determination have difficulties in producing 'significantly' better forecasts than a simple random walk model. This finding is known as the 'Meese and Rogoff Puzzle' in the exchange rate literature. Although there exists considerable disagreement about exchange rate predictability in the literature (see Rossi, 2013 for a discussion), empirical evidence seems to suggest that traditional economic predictors, such as interest rate, price and monetary differentials have rather weak predictive ability, especially when used in linear forecasting models and when considering short forecast horizons. Rossi (2013) documents that the predictive ability of fundamental variables varies noticeably across currency pairs, the type of models that are used, and over the various sample periods that are considered in the literature.

In this study, we take an entirely different modelling approach to the mainstream exchange rate forecasting literature. First, we construct an empirical heterogeneous agent model consisting of fundamentalist and chartist agents to form a model combined exchange rate forecast. More specifically, we use a simple model averaging approach to optimally weight the forecasts from two individual agent types using a 'combination regression'. Second, we use daily data to compute the agent specific forecasts. For chartist agents, a daily sampling frequency is the most natural one to adopt, as commonly used technical indicators such as moving average, momentum and relative strength index rules rely on daily data. For fundamentalist agents, which construct their forecasts based on macroeconomic information, we use daily financial variables as 'proxy variables' to obtain information about the state of the macroeconomy as it is perceived by financial market participants. The financial proxy variables that we use are yield curve data, stock price data, and data related to risk aversion and global trade activity. Third, we construct individual agent specific forecasts using the recently proposed Dynamic Model Averaging (henceforth DMA) framework. The DMA framework is an extremely flexible modelling approach, as it combines time varying parameters and model averaging into one unified framework. Forecasting agents are known to switch and adapt their prediction models over time. In order to capture this stylised fact, we employ an econometric methodology that is flexible enough to mimic this behaviour.

So far, the majority of studies in the exchange rate forecasting literature using standard macroeconomic fundamentals such as output, inflation rates, interest rates, etc., as predictors have relied upon low frequency data measured at monthly or quarterly intervals.¹ There are two major weaknesses with using standard macroeconomic data for real time forecasts of exchange rates: (i) substantial data revisions, and (ii) considerable release lags in the data. These make a 'real time' implementation and evaluation of exchange rate forecasts infeasible. The advantage of using daily financial data as a proxy for macroeconomic fundamentals is that we are able to avoid any ambiguities with respect to the timing of data releases (or its real time availability) and the impact of data revisions on the results of our study. The use of daily financial data provides a clear time stamp on what data were available to the forecasting agent at the time the forecasts were constructed.

Our empirical heterogeneous agent model allows us to include various technical indicators in the chartist predictor set in addition to macroeconomic fundamentals. The use of technical indicators as predictor variables has received considerably less attention in the recent empirical exchange rate modelling literature, despite their well documented and widespread use among practitioners.² 'Chartists', that is, agents that use technical indicators as trading signals, play also a crucial role in the theoretical finance literature on non-linear (dynamic) heterogeneous agent models (see, for instance, the classic

¹ See, for instance, Meese and Rogoff (1983), Mark (1995), Cheung et al. (2003), Wright (2008) for a few classic studies, and more recently, Della Corte et al. (2009), Della Corte and Tsiakas (2012), Della Corte et al. (2012), and Li et al. (2015) (see also Rossi, 2013 for an up-to-date major review of the exchange rate modelling literature). Notable exceptions to this mainstream literature are the studies by Rime et al. (2010), Della Corte et al. (2011) and Menkhoff et al. (2013).

² See, for instance, Allen and Taylor (1990), Taylor and Allen (1992), Menkhoff (1998), and Lui and Mole (1999) among many others.

papers by Brock and Hommes, 1997, 1998 and also more recently by De Grauwe and Grimaldi, 2005, 2006). Moreover, using a model combination approach to average the individual agent specific predictions has two advantages over simply combining the fundamentalist and chartist regressors into one large joint predictor set. First, it substantially reduces the computational burden of the construction of the DMA forecasts, which requires the computation of all possible (linear) model combinations at each point in time. Second, it allows us to provide insights about the time varying ‘importance’ of each agent’s forecast in the model combined predictions.

We use daily data for six of the most frequently traded currencies to assess the out-of-sample forecast performance of our proposed empirical heterogeneous agent model. Covering an out-of-sample evaluation period from November 2001 to June 2014, we show that forecasts from our heterogeneous agent model significantly outperform forecasts from a random walk benchmark model for all 6 currencies that we consider. More specifically, the Campbell and Thompson (2008) out-of-sample R^2 values corresponding to daily one-step ahead forecasts can be as high as 1.41%, 1.07%, 0.99%, and 0.74% for the Swiss Franc, the Euro, the Pound and the Yen series, and are somewhat lower for the Australian and Canadian Dollars at 0.29% and 0.24%. Additionally, standard statistical tests show that these forecast improvements are significant at the 10% level for the Australian and Canadian Dollars, and at the 1% level for the Euro, Yen, Pound and Swiss Franc. Some forecast gains remain statistically significant for horizons up to three days ahead, with out-of-sample R^2 values as high as 0.34%. Using a dynamic asset allocation strategy we show further that the forecasts from our combined heterogeneous agent model produce relevant ‘economic value’, yielding (annualized) out-of-sample Sharpe ratios of up to 0.89 and performance fees in excess of 460 basis points relative to a random walk benchmark.

Using various visualisation techniques, we show also that there are instabilities in the forecast performance of the model combined forecasts over the out-of-sample period that we consider. This is visible from the time series plots of the cumulative difference of the squared forecast errors, as well as the cumulative portfolio wealth, both measured relative to a random walk model’s forecasts. First, there is a consistent improvement in performance from the beginning of the out-of-sample period to the time of the Lehman Brothers collapse in September 2008. Then, from September 2008 to about February 2009, the predictive performance of the model combined heterogeneous agent forecasts increased considerably. This result is evident for all currencies from the substantial upward movement in the cumulative measures, and is contrary to the findings in Adrian et al. (2010), Molodtsova et al. (2011), and Molodtsova and Papell (2012), who report a breakdown in predictive performance over this period. Second, during this time period, the importance of the SP500 predictor variable increased homogeneously for all 6 currencies. This can be seen from the increased inclusion probabilities of the SP500 predictor variable as well as the overall increase in the DMA combined coefficient estimates. Third, there is a slowdown or breakdown in the model’s predictive performance after February 2009. Overall, there appears to be an increasing weight placed on the chartists in the agent combined forecasts from February 2009 onwards.

The remainder of the paper is organized as follows. Section 2 describes the empirical heterogeneous agent model that we construct and the model combination approach that we adopt in the paper. Section 3 describes in detail the fundamental and technical indicator data used by the two different types of agents of interest. Section 4 presents the empirical out-of-sample forecast evaluation results, together with a discussion and assessment of the economic value the model combined forecasts produce to a dynamic mean variance investor. Section 5 concludes the paper with a summary and potential future research topics.

2. Modelling approach

This section provides a detailed description of the modelling approach that we follow. We initially describe the econometric approach that we use to build the empirical heterogeneous agent model and then provide a discussion of its links and similarities to the reduced form theoretical model of De Grauwe and Grimaldi (2005). Lastly we describe the flexible econometric framework used to construct the agent specific forecasts.

2.1. Empirical model

We use a model combination (or averaging) approach to forecast the evolution of the exchange rate. This is implemented by combining the out-of-sample forecasts from two different types of agents which use two different sets of predictor variables to form their forecasts. The two agent types are (i) *fundamentalists* and (ii) *chartists*. Fundamentalists use variables that provide information about the ‘strength’ of the underlying macroeconomy to construct forecasts of the exchange rate. Chartists, on the other hand, rely upon technical indicators (or trading rules) which give an indication of momentum and trend following behaviour in exchange rates, as well as oversold or overbought conditions. Our intention here is to mimic the empirical behaviour of heterogeneous agents in foreign exchange markets by first constructing forecasts of each individual agent type and then aggregating these two forecasts (optimally) by means of a [Granger and Ramanathan \(1984\)](#) forecast combination regression. It is well known since the seminal work of [Bates and Granger \(1969\)](#) that more accurate forecasts can be obtained by combining forecasts from several models.

The model combined forecast from the two agents’ individual forecasts are constructed as follows. Denote by y_{t+1} the variable to be predicted at time $t + 1$ (ie., for simplicity of exposition, we can just think of a one-step ahead forecast here, but this is easily generalised to any forecast horizon h), and let $\mathbb{E}_t(y_{t+1}|F)$ and $\mathbb{E}_t(y_{t+1}|C)$ denote the forecasts constructed with our fundamentalists and chartists agents’ information sets, respectively. The model combined (MC) forecast is then obtained as a weighted average of the fundamentalist and chartist predictions, that is:

$$\mathbb{E}_t(y_{t+1}|MC) = (1 - \hat{\omega}_t)\mathbb{E}_t(y_{t+1}|F) + \hat{\omega}_t\mathbb{E}_t(y_{t+1}|C), \tag{1}$$

where $\hat{\omega}_t$ is the fitted value from a constrained least squares regression of the form:

$$y_t = (1 - \omega_t)\mathbb{E}_{t-1}(y_t|F) + \omega_t\mathbb{E}_{t-1}(y_t|C) + v_t, \tag{2}$$

with $\mathbb{E}_{t-1}(y_t|F)$ and $\mathbb{E}_{t-1}(y_t|C)$ being respectively the forecasts of y_t using information up to time $t - 1$, for the two agents, and v_t is an error term. Note here, that since we are using the time t estimate of ω_t to forecast y_{t+1} , this effectively implies that we assume a random walk evolution for ω_t . Also, in the empirical implementation of the forecast combination, we re-estimate ω_t for each new observation that becomes available using a rolling window scheme, so that a time-varying weight $\hat{\omega}_t$ is used in the forecast combination. Using a rolling window does not only give a more accurate representation of the real time forecasting behaviour of foreign exchange agents, but has the added benefit of providing the time-varying weights of the optimally combined forecasts and therefore the ‘influence’ or ‘activity’ of the two agent types in the model.

2.2. Relation to theoretical heterogeneous agent literature

The model combined forecasts from our empirical heterogeneous agent specification defined in (1) can be viewed as the reduced form model of the theoretical heterogeneous agent models used widely in the behavioural finance literature (see [Brock and Hommes, 1997, 1998](#)). For instance, in the model of [De Grauwe and Grimaldi \(2005\)](#), the market expectation of the change in the exchange rate is written as a weighted average of the expectations of chartist and fundamentalists, that is, using the same notation as in [De Grauwe and Grimaldi \(2005\)](#), as (see equation 8 on page 696 and including transaction costs):

$$\mathbb{E}_t(\Delta s_{t+1}) = -\eta_{ft} \overbrace{\theta(s_t - s_t^*) \mathbb{1}(|s_t - s_t^*| < C)}^{\mathbb{E}_{f,t}(\Delta s_{t+1}) = \text{fundamentalist expectation}} + \eta_{ct} \underbrace{\beta \sum_{i=0}^T a_i \Delta s_{t-i}}_{\mathbb{E}_{c,t}(\Delta s_{t+1}) = \text{chartist expectation}} \tag{3}$$

where s_t is the log of the exchange rate, s_t^* is its fundamental value, $\mathbb{1}(\cdot)$ is an indicator variable that is equal to 1 when the statement $\mathbb{1}(|s_t - s_t^*| < C)$ is true, \mathbb{E}_t denotes the expectation operator at time t , η_{ft} and η_{ct} are, respectively, 'profit functions' of fundamentalist and chartist agents that are normalised to add up to unity (see equations 4 and 5 on page 694 in De Grauwe and Grimaldi, 2005). These 'profit functions' gauge the goodness of the individual fundamentalist and chartist forecasting rules over time. From the relation in (3) it is evident that fundamentalists base their forecasts on deviations of the exchange rate from its fundamental value (ie., an 'error correction mechanism'), while chartists use a simple 'positive feedback rule' (or momentum); that is, they extrapolate past movements of exchange rate changes into the future.

Comparing our empirical model combined forecasting rule in (1) to the one proposed by the theoretical model of De Grauwe and Grimaldi (2005) in (3) highlights the similarities between the two. Both are constructed as weighted averages from the individual agents' forecasts, with the combination weights being determined by a 'goodness of performance' measure. Nevertheless, since the forecasting environment in real time is likely to be quite complex in the sense that financial market participants tend to use a variety of fundamentalist and chartist prediction rules in practice, which are also likely to change over time, we abstract from using a single fundamentalist and/or chartist predictor variable when constructing the individual agent specific predictions and use a flexible dynamic model averaging approach instead.

2.3. Constructing fundamentalist and chartist forecasts

In order to construct the model combined forecast from the two different agent types in (1), individual fundamentalist and chartist forecasts, that is, $\mathbb{E}_t(y_{t+1}|F)$ and $\mathbb{E}_t(y_{t+1}|C)$, are needed. To be able to mimic the individual forecasts from our two agents as accurately and as flexibly as possible, we use the recently proposed Dynamic Model Averaging (DMA) framework. What makes the DMA framework particularly appealing in the given context is its combination of time varying parameters and model averaging into one unifying framework, therefore mimicking the behaviour of agent learning and updating over time. We briefly describe the DMA framework in a generic setting without reference to a specific agent type and provide a more detailed description in the Appendix.

The equations of the DMA framework are:

$$y_{t+1} = \mathbf{x}_t^{(m)} \boldsymbol{\beta}_{t+1}^{(m)} + u_{t+1}^{(m)} \tag{4a}$$

$$\boldsymbol{\beta}_{t+1}^{(m)} = \boldsymbol{\beta}_t^{(m)} + \boldsymbol{\varepsilon}_{t+1}^{(m)}, \tag{4b}$$

where $m = 1, \dots, M$ denotes the model index, y_{t+1} is the variable of interest and $\mathbf{x}_t^{(m)}$ is a set of predictors included in model m . The disturbance terms $u_{t+1}^{(m)}$ and $\boldsymbol{\varepsilon}_{t+1}^{(m)}$ in (4) are jointly Multivariate Normal (MN) distributed and are mutually and serially uncorrelated, with variance and covariance $H_t^{(m)}$ and $\mathbf{Q}_{t+1}^{(m)}$. To avoid having to estimate $H_{t+1}^{(m)}$ and $\mathbf{Q}_{t+1}^{(m)}$, $H_{t+1}^{(m)}$ is modelled as an exponentially weighted moving average (EWMA) and $\mathbf{Q}_{t+1}^{(m)}$ is approximated by $(\lambda^{-1} - 1)\mathbf{P}_{t|t}^{(m)}$, with $\lambda \in [0, 1]$ acting as a 'forgetting factor' (see Raftery et al., 2010 and the Kalman Filter recursions in the Appendix for further details).

The relations in (4) show that, for each model m , we have a time-varying parameter representation for y_{t+1} . Model averaging in the DMA framework is achieved by weighting the model specific forecasts $\mathbb{E}_t(y_{t+1}^{(m)})$ by their respective predictive model probabilities which we denote by $\pi_{t|t}^{(m)}$. This quantity is computed as:

$$\pi_{t+1|t}^{(m)} = \frac{\pi_{t|t}^{\alpha(m)}}{\sum_{j=1}^M \pi_{t|t}^{\alpha(j)}}, \quad \text{with} \quad \pi_{t|t}^{\alpha(m)} = \frac{\pi_{t|t-1}^{(m)} f_N^{(m)}(y_t | \mathcal{I}_{t-1})}{\sum_{j=1}^M \pi_{t|t-1}^{(j)} f_N^{(j)}(y_t | \mathcal{I}_{t-1})}. \tag{5}$$

The term $f_N^{(m)}(y_t | \mathcal{I}_{t-1})$ in (5) is the predictive density of y_t given model m and information up to time $t - 1$, and $\alpha \in [0, 1]$ is a 'model probability' forgetting factor.

3. Data

In our empirical analysis, we use spot rates of the 6 most frequently traded currency pairs.³ We follow standard convention in the exchange rate literature and take the US Dollar to be the home currency, so that all foreign currencies are priced in US Dollars, i.e., as the US Dollar price of 1 foreign currency unit. The 6 foreign currency spot rates (denoted by S_i) are: the Euro (EUR), the Japanese Yen (JPY), the British Pound (GBP), the Australian Dollar (AUD), the Canadian Dollar (CAD) and the Swiss Franc (CHF). All exchange rate data were obtained from Bloomberg. Note that we use the 5:00 pm snap New York time as our ‘closing’ price for the exchange rates to avoid any ambiguities related to what information was still available after the closing prices were recorded.⁴ Our full data set consists of 4041 daily observations from the 4th of January 1999 to the 30th of June 2014. We use a 5 day working week in our analysis. Any missing data points due to, for instance, public holiday closings, are replaced by observations from the next previous available time period. This mimics the effective information flow as it is perceived by forecasting agents. We do not use interpolation methods to maintain the ‘real time’ aspect of the forecasts. For Japan and Switzerland, yield curve data are only available until 21st of October 2013.

We capture the influence of chartist and fundamentalist agent behaviour on exchange rates by using two different sets of predictor variables: these are (i) ‘fundamental’ variables and (ii) ‘technical’ indicators. Fundamental variables are predictors that come from standard macroeconomic models of exchange rate determination, and include measures of aggregate output, inflation and interest rates. Technical indicators are solely made up of the exchange rates’ own past values. Since we are primarily interested in a ‘real time’ (high frequency) forecast construction and evaluation, we use daily financial data as ‘proxy variables’ for fundamentals, instead of traditional (low frequency) macroeconomic variables. Our intention here is to provide as closely as possible a ‘real time’ forecast construction and evaluation scenario, which is not possible when standard low frequency data observed at monthly or quarterly intervals are used. Two major drawbacks when using monthly (or quarterly) macroeconomic variables from aggregate accounting data are that these are released with a delay and are further subject to (potentially substantial) revisions over time as new index construction methods become available and are implemented.⁵

To avoid ambiguities with respect to data releases and revisions in our forecast evaluation, we ‘extract’ information about the state of the macroeconomy from financial data, using information contained in the yield curve, stock price indices, the VIX, the TED spread, gold prices, the Baltic Dry Index (simply BDI henceforth), and the price of oil. Using financial data as a proxy for information related to macroeconomic fundamentals has the benefit of providing a clear time stamp with regard to what information was available to forecasting agents in real time (see also Harvey (1989) or Harvey (1993) for examples of other studies that use financial data as proxies for macroeconomic variables). We use simple returns in all our return calculations rather than log-returns. To conserve space, we describe the construction of all indicators, together with summary statistics of all data used in the paper in detail in the Appendix.

³ See page 11 in BIS (2013), which gives a list of the most heavily traded currency pairs by turnover. All data that we use are available from: www.danielbuncic.com/data/fx3data.zip.

⁴ We use the PX_LAST entry under the Bloomberg heading, where the price data were set manually to the New York exchange values. This is important to point out, as frequently the default setting in Bloomberg is the London close. When daily data are used, there will evidently be an overlap with the information flow generated in the US and captured by the movements in the SP500, which will then carry over into the next days’ closing price in London. Using the 5:00 pm snap New York time thus ensures that all markets have already closed on the day when the last price for the exchange rates is collected, so that this is not an issue in our analysis.

⁵ Generally, data based on aggregate accounting measures are released in intervals as new information becomes available, leading to initial, second and then final estimates. Final releases can, therefore, be up to 2–3 months after the quarter that the data are officially recorded. For details on data revisions and a standard time line of initial, second and final releases of US GDP figures, see Croushore and Stark (2001).

3.1. Fundamental variables

We use three groups of financial variables as fundamental proxies to obtain information about the state of the economy – or at least as it is perceived or expected by financial market participants. These groups are: (i) yield curve data, (ii) stock price data, and (iii) data related to risk aversion and global trade activity.⁶ From the yield curve data, we construct ‘empirical’ level (\mathcal{L}), slope (\mathcal{S}) and curvature (\mathcal{C}) factors for each foreign country, relative to the US. These are collected in

$$\mathbf{x}_t^{\text{LSC},i} = [(\mathcal{L}_t^{\text{US}} - \mathcal{L}_t^i), (\mathcal{S}_t^{\text{US}} - \mathcal{S}_t^i), (\mathcal{C}_t^{\text{US}} - \mathcal{C}_t^i)], \quad (6)$$

where i denotes the corresponding foreign country index. The stock (or equity) price based predictor set is

$$\mathbf{x}_t^{\text{EQT},i} = [r_t^{\text{EQT,US}}, r_t^{\text{EQT},i}], \quad (7)$$

Where $r_t^{\text{EQT,US}}$ and $r_t^{\text{EQT},i}$ denote the return on the SP500 and the i th foreign equity market. The risk aversion and global trade activity predictor set is

$$\mathbf{x}_t^{\text{RISK/ACTIV}} = [\Delta\text{VIX}_t, \Delta\text{TED}_t, r_t^{\text{GOLD}}, r_t^{\text{BDI}}, r_t^{\text{OIL}}], \quad (8)$$

where ΔVIX_t , and ΔTED_t are first differences of the CBOE Volatility Index and the TED spread, and r_t^i , $\forall i = \{\text{Gold}, \text{BDI}, \text{Oil}\}$ are gold, Baltic Dry Index and oil returns.

3.2. Technical variables

To enhance our ‘macroeconomic fundamentals’ information set, we construct various technical indicators. Technical analysis is a widespread method employed by market participants to forecast largely short term movements in asset prices (see also [Sarno and Taylor, 2003](#) for more details on technical trading rules and its use in foreign exchange markets). We follow the approach in [Neely et al. \(2014\)](#) and construct various technical indicators to be used as predictors in our chartist model. These technical indicators are grouped into the following 3 blocks: (i) Moving Average indicators (IMA), (ii) Momentum indicators (IMOM), and (iii) indicators based on the Relative Strength Index (RSI). The first two indicators are the same as in [Neely et al. \(2014\)](#). We add the RSI indicator, as it is another widely used technical indicator that measures the level of ‘overreaction’, that is, overbought or oversold conditions, in asset prices.

4. Forecast construction and evaluation

We now describe in detail how the forecasts are constructed and how the evaluation is carried out. Since we are primarily interested in the real time predictive performance of the model, we implement an out-of-sample forecast evaluation of the model. Also, before we outline in detail the statistical criteria that we utilise to assess the performance of the model, and before the results of the forecast evaluation are reported, we initially describe the prediction setting that we use in our evaluation.

4.1. Prediction setting

As outlined in [Section 2](#), we implement an agent based model averaging/combination approach to forecast the returns of our 6 exchange rates of interest. The agent based model combined predictions (henceforth, simply MC predictions, which we denote by $\hat{r}_{t+h|t}^{\text{MC}}$) are computed as:

⁶ See the Appendix for a more elaborate motivation of using these groups of variables as fundamental predictors.

$$\hat{r}_{t+\eta|t}^{MC} = (1 - \hat{\omega}_t) \hat{r}_{t+\eta|t}^F + \hat{\omega}_t \hat{r}_{t+\eta|t}^C, \tag{9}$$

where $\hat{r}_{t+\eta|t}^F$ and $\hat{r}_{t+\eta|t}^C$ denote, respectively, the individual 1-step ahead forecasts constructed by fundamentalist and chartist agents at time t . The (fitted) weights $\hat{\omega}_t$ in (9) are obtained by means of a Granger and Ramanathan (1984) forecast combination regression of the form:

$$r_t = (1 - \omega_t) \hat{r}_{t|t-1}^F + \omega_t \hat{r}_{t|t-1}^C + v_t \tag{10}$$

which, for estimation purposes, can be conveniently re-written in the equivalent form:

$$(r_t - \hat{r}_{t|t-1}^F) = \omega_t (\hat{r}_{t|t-1}^C - \hat{r}_{t|t-1}^F) + v_t, \tag{11}$$

where the weights ω_t are restricted to be in the $[0, 1]$ interval.

Two features are evident from the specification of the model averaged predictions defined in (9) and how the combination weights are obtained. First, the time t estimate of ω_t is used in the averaged forecast construction in (9). This implies that the best forecast of the averaging weight ω_t one period into the future is its time t estimate, that is, $\mathbb{E}_t(\omega_{t+1}) = \hat{\omega}_t$. The weights are thus assumed to evolve as a random walk process. Second, individual return forecasts from the fundamentalist and chartist agents are needed. We describe below in the next section how these are computed using the DMA methodology that was outlined in Section 2.3.

4.1.1. Computing individual fundamentalist and chartist forecast

Let $A = \{F, C\}$ denote the agent type for which we construct the forecast, i.e., a fundamentalist or chartist agent. Following the general description of the DMA framework in Section 2.3, the return forecast of agent A is computed as:

$$\hat{r}_{t+\eta|t}^A = \sum_{m=1}^M \mathbf{x}_t^{A(m)} \hat{\beta}_{t|t}^{A(m)} \pi_{t+\eta|t}^{A(m)}, \tag{12}$$

where the fundamentalist and chartist predictor sets are

$$\begin{aligned} \mathbf{x}_t^F &= [1, \mathbf{x}_t^{LSC}, \mathbf{x}_t^{EQT}, \mathbf{x}_t^{RISK/ACTIV}] \quad \text{and} \\ \mathbf{x}_t^C &= [1, r_t, \text{IMA}_t^{(200)}, \text{IMOM}_t^{(130)}, \text{RS}_t^{(14)}], \end{aligned}$$

respectively.⁷ The sequence $\hat{\beta}_{t|t}^{A(m)}$ are the time t filtered estimates of $\beta_t^{A(m)}$, which are obtained by standard Kalman Filter recursions (see again the Appendix for details). The last term in (12) are forecasts of the model probabilities as defined in (5). The smoothing term in the EWMA specification of $H_t^{(m)}$ in (4) is fixed at 0.94. The α and λ parameters are ‘calibrated’ at 0.99 and 0.999, respectively.⁸

⁷ The 200, 130 and 14 superscripts on the technical indicators denote the number of trading days they are formed over (see the Appendix for further details). In \mathbf{x}_t^F and \mathbf{x}_t^C , the intercept term is denoted by 1, and the exchange rate return r_t in \mathbf{x}_t^C allows for the possibility of AR type dynamics in the returns. Also, it should be clear that there is only very weak evidence of any autocorrelation in the return series from the summary statistics that we report in the Appendix. Nevertheless, these summary statistics are an unconditional measure computed over the full sample period. There may exist times when an increase in the autocorrelation of returns (in absolute terms) occurs. To allow for this possibility, we include a one period lagged exchange rate return in the chartist predictor set. This formulation is also in line with the dynamics specified in De Grauwe and Grimaldi (2005) for chartist agents (see their equation 3 on page 694).

⁸ Exact details on how this calibration was performed are provided in the Appendix. Also, to implement the Kalman Filter, we need to specify initial values. We follow Koop and Korobilis (2012) and use a diffuse prior for $\hat{\beta}_{0|0}^{A(m)} \sim \text{MN}(\mathbf{0}_{K_m}, 100\mathbf{I}_{K_m})$, where $\mathbf{0}_{K_m}$ is a $(K_m \times 1)$ dimensional vector of zeros and \mathbf{I}_{K_m} is $(K_m \times K_m)$ dimensional identity matrix. The model updating probabilities for agent type A are initialised with an uninformative prior $\pi_{0|0}^{(m)} = \frac{1}{M}$, so that all models are assumed to be equally likely.

4.1.2. Fitting and evaluation periods

Our entire available data set consists of $T = 4039$ ($T = 3860$) observations, covering the period from January 4, 1999 to June 30, 2014 for the EUR, GBP, AUD and CAD, and for JPY and CHF from January 4, 1999 to October 22, 2013. Since we are primarily interested in an out-of-sample forecast evaluation of the combined or model averaged predictions from the two different forecasting agents, we effectively require two *in-sample* or *fitting* periods. That is, we need one *calibration/initialisation* period to obtain the agent specific forecasts, $\hat{r}_{t+1|t}^F$ and $\hat{r}_{t+1|t}^C$, and another *calibration* period for the combination weights $\hat{\omega}_t$ in (9) to be determined. We use the first 500 observations (from January 6, 1999 to December 5, 2000) as the *in-sample* period to calibrate the λ and α parameters, which in our Kalman Filter setting also serves as a *burn-in* period to minimise the influence of the initial values (or priors) on the filtered state vector $\hat{\beta}_{t|t}$. We then take the next 252 observations from December 6, 2000 to November 22, 2001 – which corresponds to a one year time horizon – to get our estimate $\hat{\omega}_t$ needed to compute our first model averaged out-of-sample forecast for November 23, 2001. We then roll through the rest of the out-of-sample data to update $\hat{\omega}_t$ using a fixed window size of 252 observations and produce (recursively updated) one-step ahead forecasts of the returns. The effective out-of-sample period thus spans from November 23, 2001 to June 30, 2014 (to October 22, 2013 for JPY and CHF), yielding 3287 (respectively 3102) evaluation points.

4.2. Evaluation criteria

We assess the out-of-sample forecast performance of the proposed agent based averaging framework by following the recent literature on forecasting the equity premium. That is, we follow the approach of Rapach et al. (2013), Neely et al. (2014) and others and evaluate the forecasts in terms of the Campbell and Thompson (2008) out-of-sample R^2 (denoted by R_{os}^2 henceforth) and the Clark and West (2007) Mean Squared Forecast Error (MSFE) adjusted t -statistic. In all our out-of-sample forecast evaluations, we use the forecasts from a driftless random walk (RW) as the benchmark in the statistical tests.

The Campbell and Thompson (2008) R_{os}^2 is computed as:

$$R_{os}^2 = 1 - \frac{MSFE_{(MC)}}{MSFE_{(RW)}}, \tag{13}$$

where $MSFE_{(\ell)} = \frac{1}{T_{os}} \sum_{t=T_{is}}^T \hat{e}_{t+1|t}^{2(\ell)}$ and $\hat{e}_{t+1|t}^{(\ell)} = (r_{t+1} - \hat{r}_{t+1|t}^{\ell})$ the (one-step ahead) forecast errors from model ℓ , where $\ell = \{RW, MC, F, C\}$.⁹ The Clark and West (2007) MSFE adjusted t -statistic is (denoted by CW-statistic henceforth) is computed as:

$$CW - \text{statistic} = \frac{\overline{c\bar{w}}}{\sqrt{\text{Var}(\overline{c\bar{w}})}} \tag{14}$$

where $\overline{c\bar{w}} = T_{os}^{-1} \sum_{t=T_{is}}^T c\bar{w}_{t+1}$, $c\bar{w}_{t+1} = \hat{e}_{t+1|t}^{2(RW)} - \hat{e}_{t+1|t}^{2(MC)} - [\hat{r}_{t+1|t}^{(RW)} - \hat{r}_{t+1|t}^{(MC)}]^2$ and $\text{Var}(\overline{c\bar{w}})$ are the variance of the sample mean. We also compute and plot the cumulative difference of the squared forecast errors (denoted by cumSFE) of the RW and MC forecasts over the out-of-sample period.¹⁰

⁹ The terms T_{os} and T_{is} denote, respectively, the number of out-of-sample and in-sample observations, so that $T_{is} + T_{os} = T$, with T being the full sample size.

¹⁰ In our setting, this difference is computed as: $\text{cumSFE}_{t+1} = \sum_{t=T_{is}}^{T_{os}} (\hat{e}_{t+1|t}^{2(RW)} - \hat{e}_{t+1|t}^{2(MC)})$. A value of cumSFE above zero indicates that the model combined forecasts are more accurate. In general, a rising value in cumSFE means that the MC forecasts produce better predictions than the RW benchmark.

4.3. Forecast evaluation results

In this section, we provide a detailed statistical as well as economic analysis of the forecasts performance of our model. For readability, we break this section into 5 subsections, separately discussing the statistical evaluation results, the time varying evolution of the agent weight function, what is driving the predictability results, and a statistical evaluation over three sub-periods to measure the performance before, during and after the financial crisis. Lastly, we discuss the economic significance of our statistical results.

4.3.1. Statistical forecast evaluation

In [Table 1](#) we present the one-step ahead out-of-sample forecast evaluation results for the period from November 23, 2001 to June 30, 2014 (to October 22, 2013 for JPY and CHF). The first column in [Table 1](#) shows the models that are fitted, the second column shows the mean squared forecast errors (MSFEs), the third column the MSFEs relative to the RW benchmark, the fourth column the [Campbell and Thompson \(2008\)](#) R_{os}^2 (in percent) as defined in (13), and the fifth and sixth columns display the

Table 1
One-step-ahead out-of-sample forecast evaluation results.

Model	MSFE	Relative-MSFE	R_{os}^2 (%)	CW-statistic	p-Value
EUR					
Random walk (RW)	0.3839	–	–	–	–
Chartist (C)	0.3847	1.0021	–0.2102	–0.1293	0.5514
Fundamentalist (F)	0.3805	0.9911	0.8908	3.5743	0.0002
Model combined (MC)	0.3796	0.9888	1.1231	3.6361	0.0001
JPY					
Random walk (RW)	0.4249	–	–	–	–
Chartist (C)	0.4257	1.0019	–0.1898	–0.3047	0.6197
Fundamentalist (F)	0.4211	0.9910	0.8990	3.6487	0.0001
Model combined (MC)	0.4217	0.9925	0.7504	3.2288	0.0006
GBP					
Random walk (RW)	0.3239	–	–	–	–
Chartist (C)	0.3244	1.0016	–0.1565	–0.2143	0.5848
Fundamentalist (F)	0.3213	0.9920	0.8041	3.2922	0.0005
Model combined (MC)	0.3209	0.9907	0.9321	3.3830	0.0004
AUD					
Random walk (RW)	0.7241	–	–	–	–
Chartist (C)	0.7244	1.0004	–0.0410	0.5078	0.3058
Fundamentalist (F)	0.7249	1.0012	–0.1176	1.2770	0.1008
Model combined (MC)	0.7221	0.9972	0.2777	1.5899	0.0559
CAD					
Random walk (RW)	0.3662	–	–	–	–
Chartist (C)	0.3676	1.0037	–0.3713	–1.6913	0.9546
Fundamentalist (F)	0.3654	0.9977	0.2347	1.6412	0.0504
Model combined (MC)	0.3655	0.9979	0.2115	1.3848	0.0831
CHF					
Random walk (RW)	0.4962	–	–	–	–
Chartist (C)	0.4978	1.0031	–0.3097	–0.6766	0.7507
Fundamentalist (F)	0.4902	0.9878	1.2157	4.1641	0.0000
Model combined (MC)	0.4897	0.9868	1.3209	4.0895	0.0000

Notes: This table reports the one-step ahead out-of-sample forecast evaluation results for the 6 exchange rates of interest. In column 1, the various models that are considered in the evaluation are listed. Columns 2 and 3 show the mean squared forecast errors (MSFE) and Relative-MSFEs corresponding to the various models, where the Relative-MSFE is computed by taking the values from column 2 and deflating them by each exchange rate's respective $MSFE_{(RW)}$. In column 4 we report the [Campbell and Thompson \(2008\)](#) out-of-sample R^2 (denoted by R_{os}^2) which is computed as $1 - MSFE_{(i)} / MSFE_{(RW)}$, $\forall i = \{C, F, MC\}$. The [Clark and West \(2007\)](#) CW-statistic and its corresponding p-value are given in columns 6 and 7. The out-of-sample evaluation period is from November 23, 2001 to June 30, 2014 for all currencies, except for the JPY and CHF series, where the sample ends already in October 22, 2013.

Clark and West (2007) CW-statistic and its corresponding one-sided p -value. Note that the forecasts that are listed here are from the benchmark RW model, the individual chartist and fundamentalist forecasts, as well as the model combined forecasts, where the ‘averaging’ weights were obtained from the Granger and Ramanathan (1984) combination regressions in (11).

We can initially notice from the results in Table 1 that chartist forecasts as a whole seem to perform rather poorly when compared to fundamentalist and RW forecasts. This is evident from chartist forecasts producing the largest MSFEs. It is further evident that fundamentalist forecasts generate lower MSFEs than the benchmark RW model, producing out-of-sample R^2 values that are positive and as high as 1.22% for the Swiss Franc, with the Australian Dollar being the only currency with a negative R^2_{os} of -0.18% . From the magnitude of the CW-statistics and their p -values, one can see that improvements in forecast accuracy of the fundamentalist forecasts are statistically significant at the 1% level for the Euro, the Yen, the Pound, and the Swiss Franc and borderline insignificant at the 5% and 10% levels for the Canadian and Australian Dollars, respectively.

Although the above reported out-of-sample R^2 may appear small in magnitude and thus unimportant, we should highlight here that even seemingly low R^2_{os} values can be economically sizeable. In a broader context, one should expect to see only a very small predictive component in exchange rate returns, if foreign exchange markets are believed to be efficient. Also, our R^2_{os} are computed at the daily frequency. To put this magnitude into perspective, in the equity premium forecasting literature, Campbell and Thompson (2008), and more recently Neely et al. (2014), have shown that R^2_{os} values as low as 0.5% computed on monthly data produce economically meaningful predictive results in the sense that ‘large’ gains in portfolio performance can be obtained which an investor would be willing to pay for. Similarly, in the exchange rate forecasting literature, Della Corte and Tsiakas (2012) and Li et al. (2015) have shown that modest improvements in the out-of-sample R^2 relative to a benchmark random walk model can generate annual performance fees in the order of magnitude of more than 4% per year for a risk-averse investor. We return more formally to answer the question of economic significance of the model combined forecast improvements over the random walk ones in detail in Section 4.3.5.

Looking over the model combined forecast results shown in the last row of each currency grouping in Table 1, we see that, with the exception of the JPY series, all other currency returns benefit from the performance based weighting of the two agent’s predictions in the construction of the combined forecasts. The magnitude of this gain can be seen from the higher R^2_{os} values for the EUR, GBP, AUD, CAD and CHF series. For these 5 series, the biggest forecast gain in terms of a higher R^2_{os} is obtained for the Australian Dollar, which increased by approximately 0.36 percentage points and the lowest for the Canadian Dollar, which only improved by about 0.01 percentage points. Note here also that, although chartist forecasts are inferior to fundamentalist forecasts when averaged over the full out-of-sample period, there do exist instances where chartist predictions perform better than forecasts based solely on fundamental predictor variables. It is exactly this feature of the MC forecasts that leads to an overall improvement when averaging over the two individual agent based predictions.

To gain a better understanding of the positive (statistical) forecast evaluation results that we obtain, we examine the evolution of the cumulative difference of the MSFEs of the model combined forecasts (relative to the RW benchmark) over time. This cumSFE series is plotted in Fig. 1 for the 6 currencies of interest. Note here that, because of the shorter out-of-sample evaluation period for the Yen and the Franc, the cumSFE series ends already on October 22, 2013 for these two currencies. Nevertheless, for reasons of comparability with respect to date entries in the figures that we show, we have plotted all series up to June 30, 2014 and set a common y -axis scale over the interval $[-5, 35]$. As a reminder, the cumSFE series is defined such that an increasing value indicates an improvement in the MC predictions relative to the RW benchmark; that is, the RW benchmark produces larger one-step ahead forecast errors.

Looking over the cumSFE series plotted in Fig. 1, one can notice the following 4 visually striking features from these plots. First, the $cumSFE_t$ is (nearly) uniformly above zero for all currencies over the whole out-of-sample evaluation period, except for the AUD series, where it is above zero only up to June 2010, dropping below 0 thereafter. Second, the cumSFE series is (nearly) monotonically increasing for all series up to the September 2008 period, i.e., around the time of the Lehman Brothers collapse. Third, the effect of the Lehman Brothers collapse has a strong impact on the predictability




Fig. 1. Time series evolution of the cumulative sum of squared forecast errors (cumSFE,) of the model combined forecasts (relative to the RW model) over the out-of-sample period from November 23, 2001 to June 30, 2014 (October 22, 2013 for the CHF and JPY series). Note that we plot all series up to June 30, 2014 so that the dates on the x-axis can be compared easily.

results for all 6 currencies.¹¹ From September 2008 until approximately February 2009, the predictability in all 6 series (relative to the RW) experienced a huge boost. This effect is most evident for the Australian Dollar and the Swiss Franc, and least so for the Canadian Dollar. Fourth, since the Lehman Brothers collapse in September 2008, it has become more difficult to outperform the benchmark random walk predictions. This conclusion can be reached from the overall downward *‘trend’* that is visible in the cumSFE series from approximately February 2009 until the end of the sample period in June 2014 (October 2013).

The single most peculiar visual result of the cumSFE series is for the Australian Dollar, where the forecast accuracy of the model combined predictions initially improves dramatically, with the cumSFE increasing from a value of approximately 19 on October 9, 2008 to nearly 28 on October 10, 2008, then dropping back to 19 on October 13, and decreasing rapidly thereafter. By October 21, 2008, the

¹¹ A similar result is found in [Buncic and Moretto \(2015\)](#) in a copper forecast evaluation and also in [Buncic and Gislser \(2015\)](#) in a volatility forecasting setting.

cumSFE had dropped below 6. For the Euro and the Canadian Dollar, similar but less accentuated swings in the cumSFE are visible. For the Euro, the cumSFE increased from a value of around 12 on September 9, 2008 to nearly 20 by October 6, 2008, dropping back to just below 13 on November 21, 2008, and rising up to 21 by December 19, 2008, before gradually dropping, reaching a value of 15 by May 2009. For the Canadian dollar, although at an overall smaller cumSFE magnitude, also fairly abrupt movements are visible. For instance, the cumSFE stood at about 5 on September 29, 2008, and then doubled to over 10 by October 10, 2008. By February 2010, it had declined to a value of 2. The Yen, Pound and Franc experienced more stable and seemingly permanent increases in predictability over the September 2008 to May 2009 period. This is highlighted by the cumSFE series increasing from values of around 6 to 14, 5 to 12, and 15 to 25, for the JPY, GBP and CHF, respectively over this time span. All in all, it should be clear from the visual analysis that we presented here that around the time of the Lehman Brothers collapse a substantial boost in exchange rate predictability was realised.¹²

4.3.2. Time series evolution of the model combination weights

To learn about the influence of chartist agents, we show the time series evolution of the weight function $\hat{\omega}_t$, as defined in (9), in Fig. 2. $\hat{\omega}_t$ is plotted as a red solid line in Fig. 2. We also create an indicator variable that is equal to 1 if $\hat{\omega}_t > 0.5$. This indicator variable is drawn as a grey shaded area in Fig. 2. The time series evolution of the weight functions portray a number of interesting insights. First, we can notice that the weight of chartists agents in the model combined forecasts varies considerably between the 6 currencies of interest. The influence of chartist agents was rather low for the first four years into the out-of-sample period for the EUR, the AUD and the CHF, with some episodes of importance for GBP, CAD and JPY over the 2002 to 2003 period. However, from the mid to end of 2009 onwards, $\hat{\omega}_t$ began to increase steadily. Recall that this period coincides with the strong rebound in global equity prices that followed the bottom of the bear market in March 2009 and also the flow on effects of the implementation of quantitative easing in late November 2008 in the US. This increase in chartist weights is particularly noticeable for the EUR, the GBP and AUD series. For the Yen, chartist agents did not bear any important effect on the forecasts until the beginning of March 2011, while for the Canadian Dollar, the period following the bottom of the equity bear market in March 2009 to March 2011 was, apart from a short time interval from March 2010 to May 2010, largely driven by chartist forecasts.

4.3.3. What is driving the predictability results

Given the positive forecast performance of the model combined forecasts, we now examine the dynamic influence of the individual fundamentalist and chartist predictors over time. To do this, we examine the dynamic evolution of two quantities of interest. The first one is the posterior inclusion probability (PIP_t for short) of the predictor variables. The second one is the weighted average of the updated estimates of the latent state vector, which we denote by $\hat{\beta}_{it}^{(DMA)}$. The $(k \times 1)$ dimensional PIP vector at time t is computed as:

$$PIP_t = \sum_{m=1}^M \pi_{it}^{(m)} \mathbb{1}(\mathbf{x}_t \in \mathcal{M} = m), \quad (15)$$

where $\pi_{it}^{(m)}$ is the updated model probability and $\mathbb{1}(\mathbf{x}_t \in \mathcal{M} = m)$ is an indicator variable that is equal to 1 if any of the regressors in \mathbf{x}_t are included in the m th model. Intuitively, the PIP is constructed by simply summing over all updated model probabilities $\pi_{it}^{(m)}$ that contain the i th predictor variable $\{x_{i,t}\}_{i=1}^k$ in its model set.

¹² There are several papers that find some predictability of exchange rates before the financial crisis, with their results, nevertheless, breaking down after 2008/2009 (see, for instance, Adrian et al., 2010, Molodtsova et al., 2011, and Molodtsova and Papell, 2012, and references therein). Our improved predictability result over the crisis period is thus entirely new. These existing studies, nevertheless, use standard low frequency data, and thus face an entirely different conditioning set upon which the forecasts are based.




Fig. 2. Time series evolution of the predicted chartist weight (or influence) function $\hat{\omega}_t$ over the out-of-sample forecasting period from November 23, 2001 to June 30, 2014 (October 22, 2013 for the CHF and JPY series). The weight function is obtained from a rolling window regression with fixed estimation window of 1 year (252 daily observations). That is, we get an estimate of ω_t over the sample period from December 6, 2000 to November 22, 2001 and use this first estimate as the predicted weight for November 23, 2001 and then roll through the out-of-sample observations.

The weighted average of the time varying parameter estimates is obtained as:

$$\hat{\beta}_{t|t}^{(DMA)} = \sum_{m=1}^M \pi_{t|t}^{(m)} \hat{\beta}_{t|t}^{(m)} \mathbb{1}(\mathbf{x}_t \in \mathcal{M} = m), \tag{16}$$

where $\hat{\beta}_{t|t}^{(m)}$ are again the time t filtered estimates of $\beta_t^{(m)}$ from the Kalman Filter recursions. Thus, $\hat{\beta}_{t|t}^{(DMA)}$ is computed by averaging each model's $\hat{\beta}_{t|t}^{(m)}$ estimate over the updated model probabilities $\pi_{t|t}^{(m)}$ at each point in time for those regressors that are included in the model, with the averaging weight determined by $\pi_{t|t}^{(m)}$. These two quantities are plotted in Figs. 3–8 for all fundamentalist and chartist predictor variables.

Each plot in Figs. 3–8 shows the $\hat{\beta}_{t|t}^{(DMA)}$ coefficients marked by the blue solid line (left-axis scale) and the PIP, drawn as the red solid line (right-axis scale) for all predictor variables over the entire

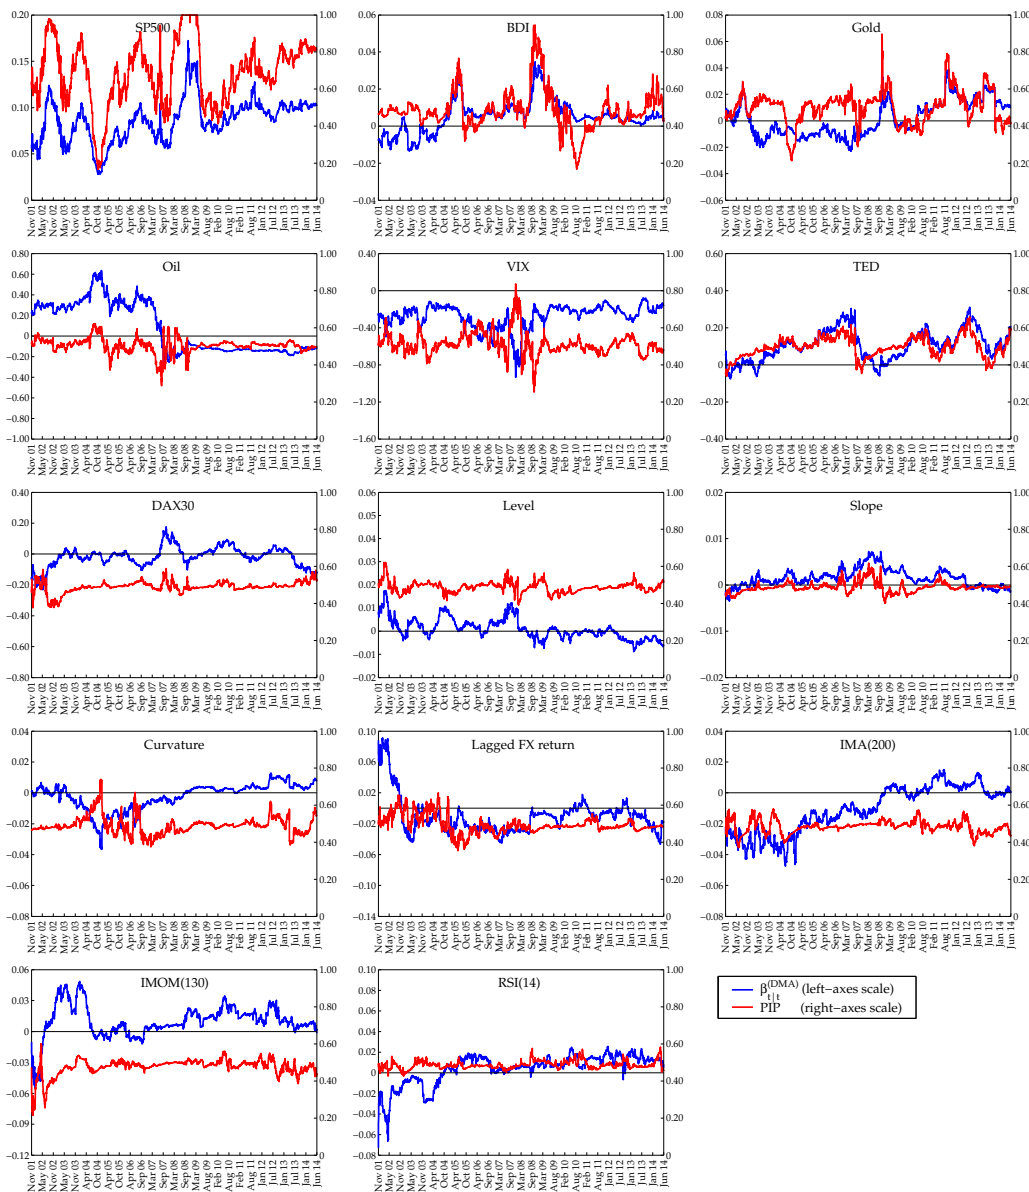


Fig. 3. Time series evolution of the (DMA) model averaged $\hat{\beta}_{it}$ estimates (blue solid line, left-axis scale) and the posterior inclusion probabilities (PIPs) (red solid line, right-axis scale) for the EUR series over the out-of-sample forecast evaluation period from November 23, 2001 to June 30, 2014. The plots show the parameter estimates of all the fundamentalist (plots 1 to 10) as well as the chartist predictor variables (plots 11 to 14), respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

out-of-sample period from November 23, 2001 to June 30, 2014. We intentionally set the same axes scales in all 6 plots to be able to easily compare the magnitudes of the PIP_t and $\hat{\beta}_{it}^{(DMA)}$ sequences across the exchange rates. There are a number of interesting features visible from these figures, which can be summarized as follows. First, the SP500 has a positive coefficient for all 6 exchange rates that we

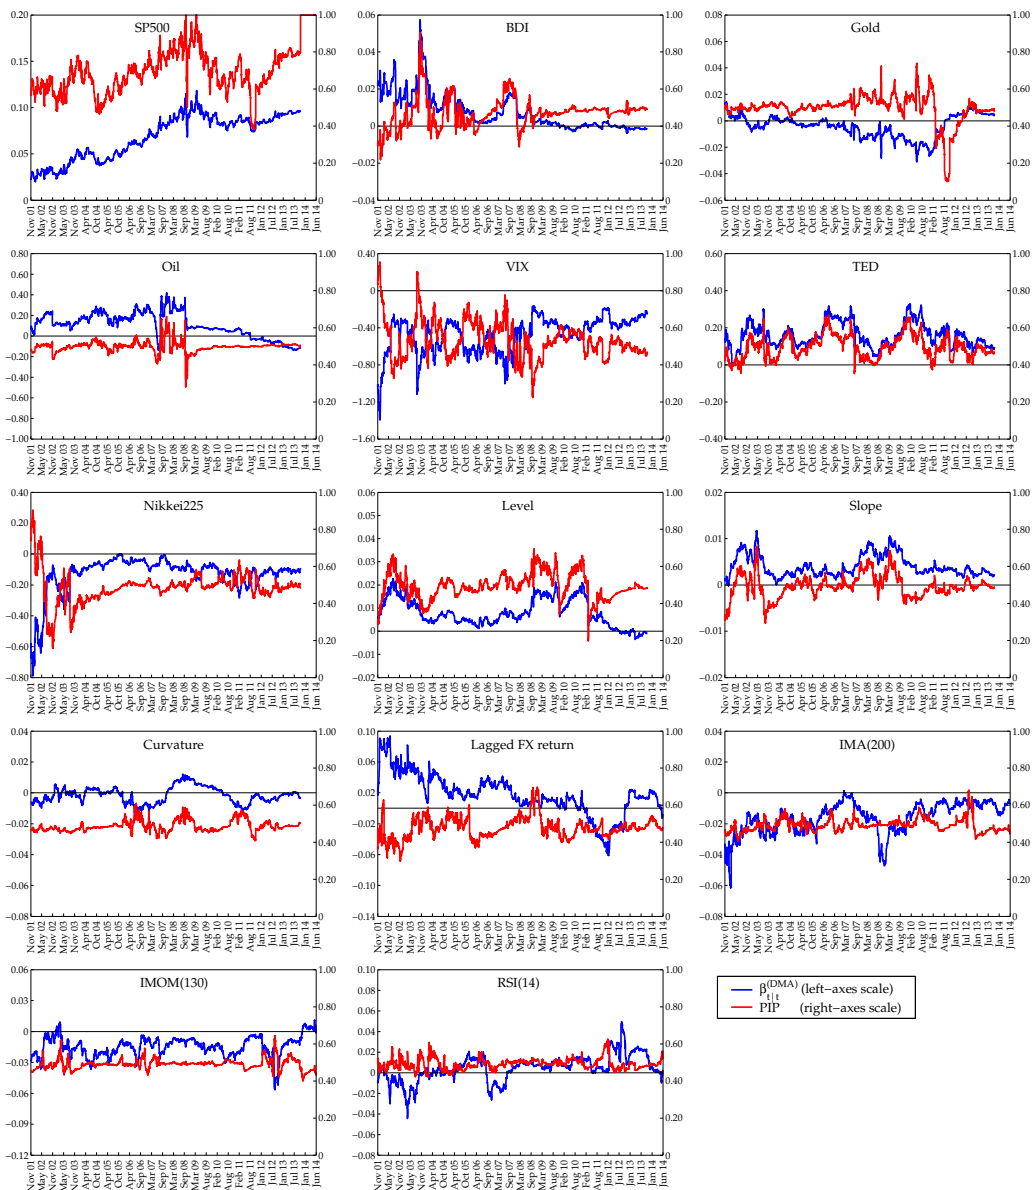


Fig. 4. Time series evolution of the (DMA) model averaged $\hat{\beta}_t$ estimates (blue solid line, left-axis scale) and the posterior inclusion probabilities (PIPs) (red solid line, right-axis scale) for the JPY series over the out-of-sample forecast evaluation period from November 23, 2001 to June 30, 2014. The plots show the parameter estimates of all the fundamentalist (plots 1 to 10) as well as the chartist predictor variables (plots 11 to 14), respectively. The fundamentalist based forecasts end on October 22, 2013 (due to the yield curve data). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

consider. What is striking here is the mild upward trending evolution in not only the model averaged parameter estimates, but also in the posterior inclusion probabilities. Although there are some differences in the magnitudes of the PIP_t and $\hat{\beta}_{it}^{(DMA)}$ series during the time of the Lehman Brothers collapse across the 6 series, all indicate that the importance of the information contained in the SP500

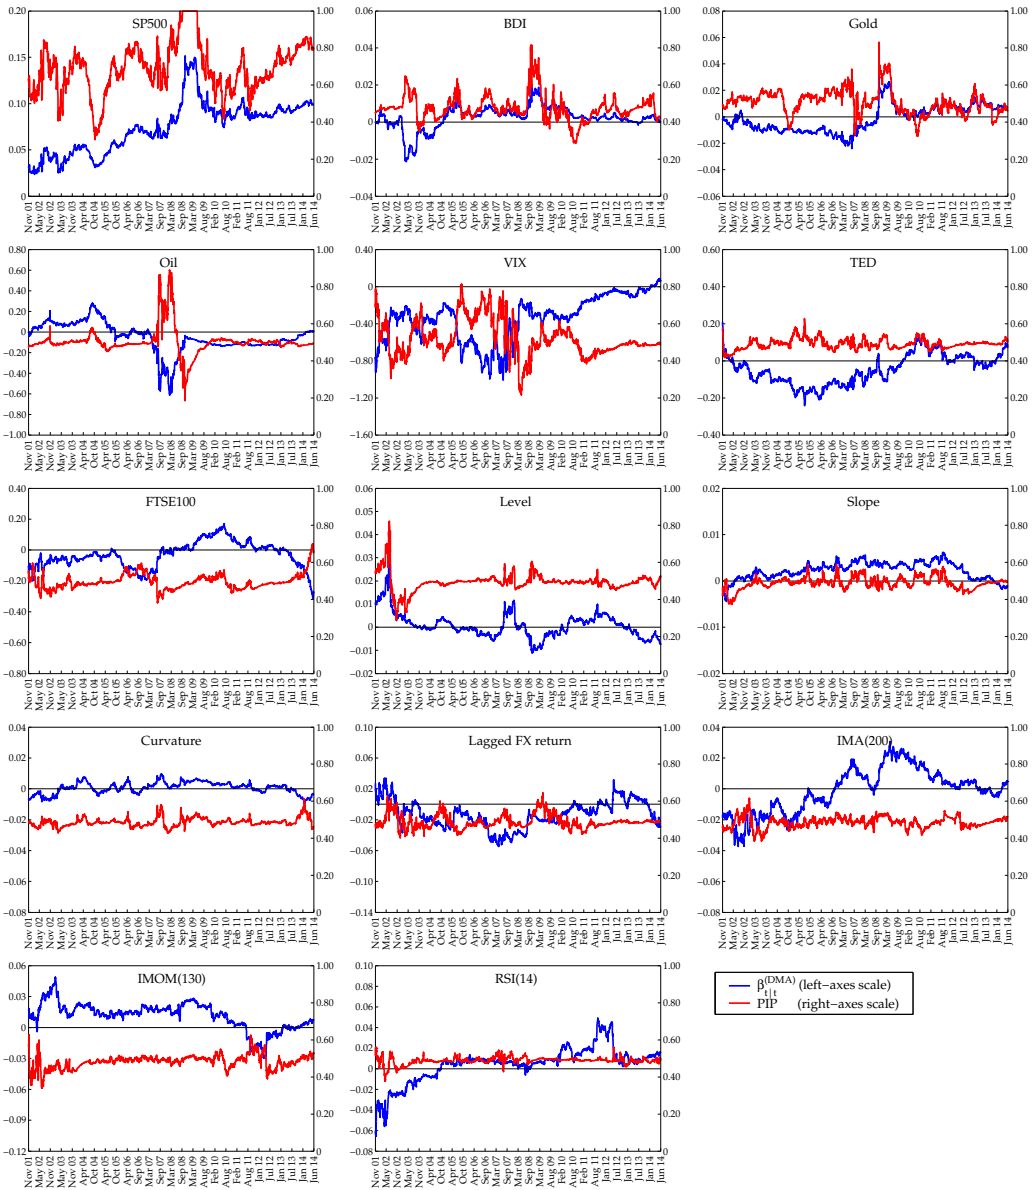


Fig. 5. Time series evolution of the (DMA) model averaged $\hat{\beta}_{it}$ estimates (blue solid line, left-axis scale) and the posterior inclusion probabilities (PIPs) (red solid line, right-axis scale) for the GBP series over the out-of-sample forecast evaluation period from November 23, 2001 to June 30, 2014. The plots show the parameter estimates of all the fundamentalist (plots 1 to 10) as well as the chartist predictor variables (plots 11 to 14), respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

predictor variable increased substantially from September 2008 until about February 2009, dropping back towards its 'trend level' thereafter. The PIP_t for the SP500 is equal to (or very close to 1) over this period. This means that the best forecasting models, i.e., the ones with the highest predictive probability, always include the SP500 as a predictor variable.

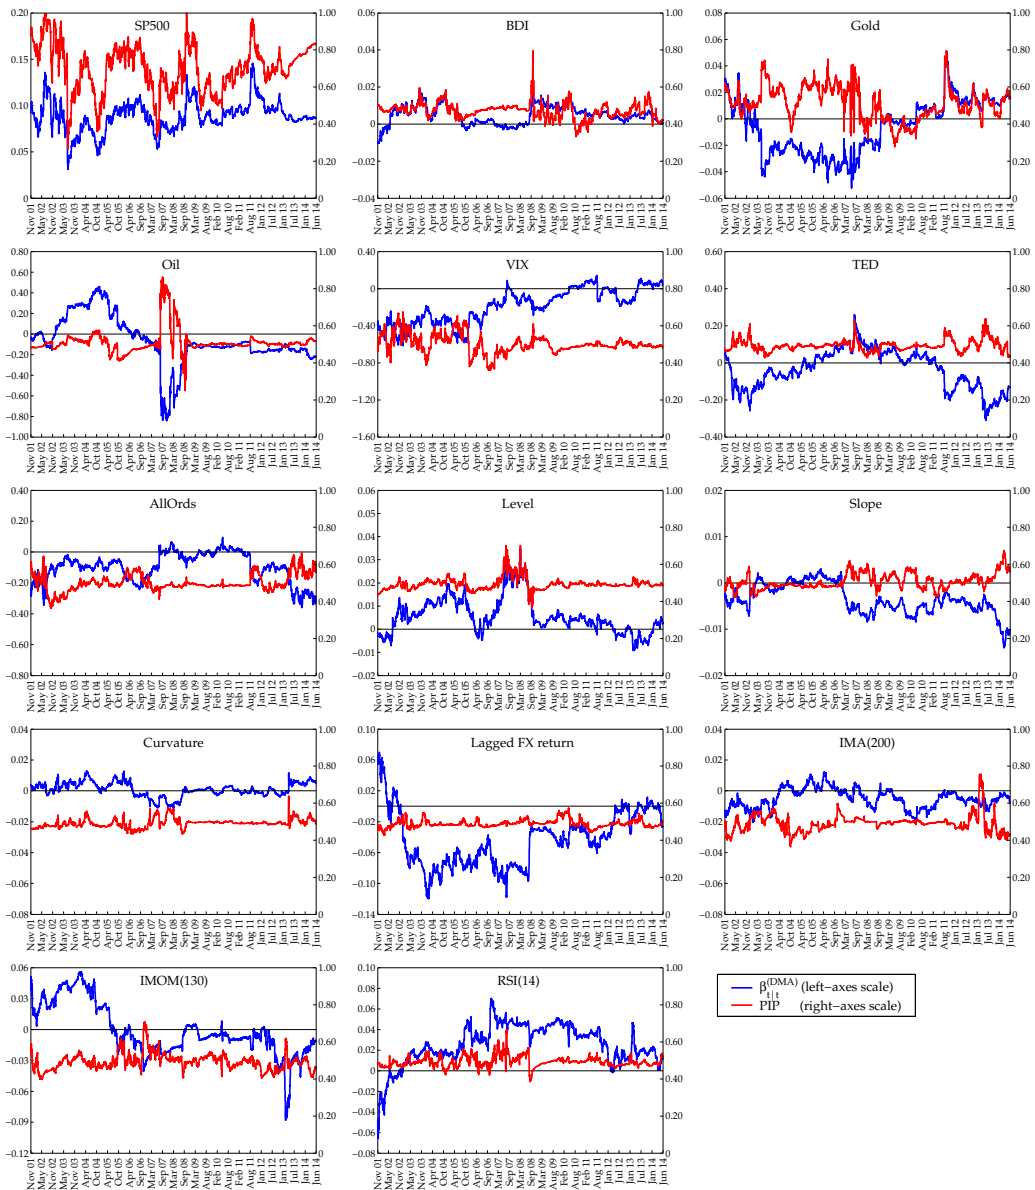


Fig. 6. Time series evolution of the (DMA) model averaged $\hat{\beta}_t$ estimates (blue solid line, left-axis scale) and the posterior inclusion probabilities (PIPs) (red solid line, right-axis scale) for the AUD series over the out-of-sample forecast evaluation period from November 23, 2001 to June 30, 2014. The plots show the parameter estimates of all the fundamentalist (plots 1 to 10) as well as the chartist predictor variables (plots 11 to 14), respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Second, the influence of the Gold, Oil and BDI series increased also somewhat during this period for some of the currencies. What is rather surprising to see is that the influence of the VIX and the TED spread is not as sizeable during the time of the Lehman Brothers collapse as reported in other studies (ie., Brunnermeier et al., 2009). For instance, for the AUD series, the magnitude of the TED spread

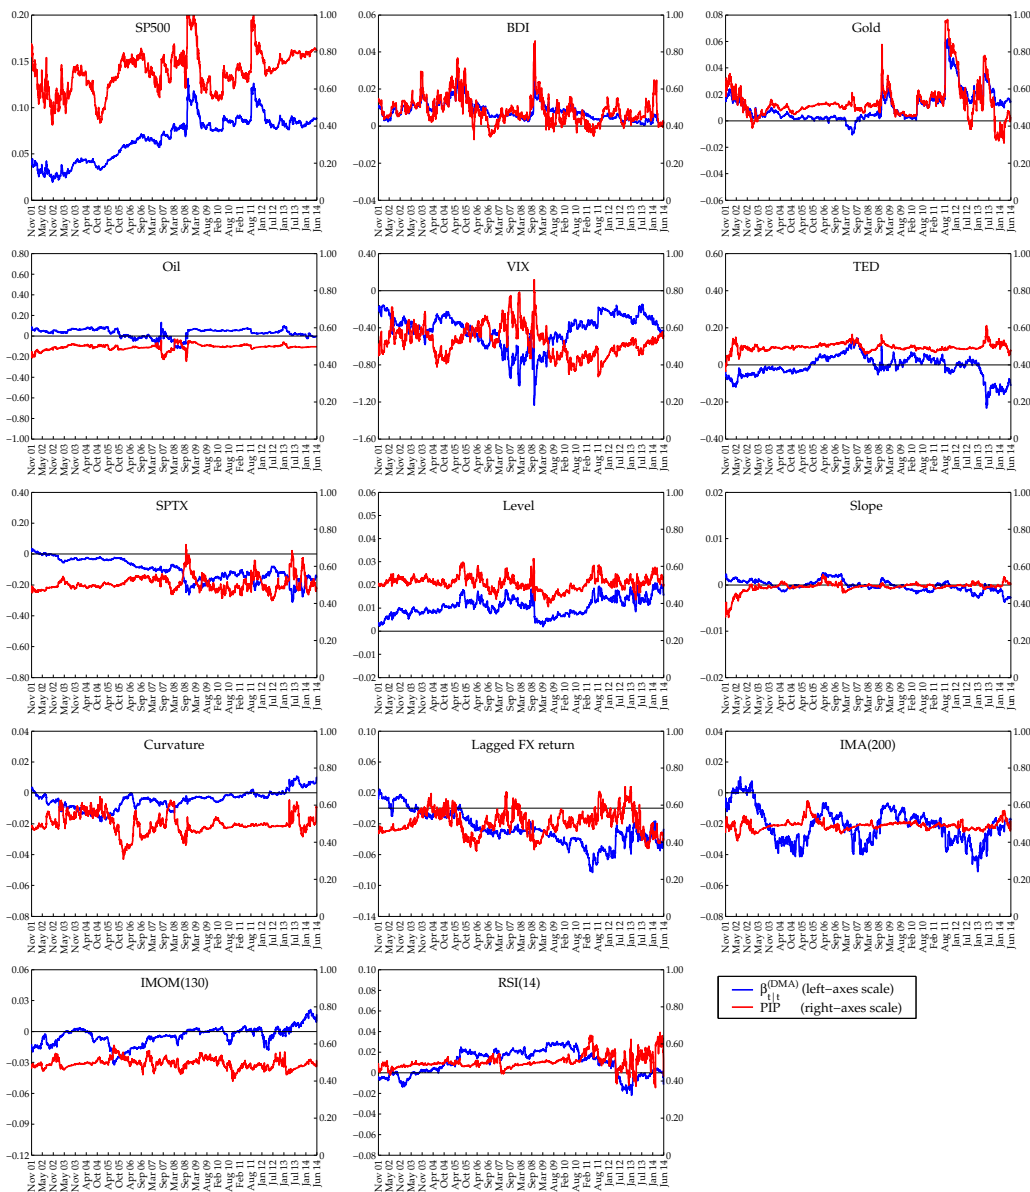


Fig. 7. Time series evolution of the (DMA) model averaged $\hat{\beta}_t$ estimates (blue solid line, left-axis scale) and the posterior inclusion probabilities (PIPs) (red solid line, right-axis scale) for the CAD series over the out-of-sample forecast evaluation period from November 23, 2001 to June 30, 2014. The plots show the parameter estimates of all the fundamentalist (plots 1 to 10) as well as the chartist predictor variables (plots 11 to 14), respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

coefficient jumped up (increase in magnitude) around the August/September 2007 period, when first signs of liquidity shortages were realised and interbank lending slowed down, while the magnitude of the coefficient dropped towards zero in September 2008, suggesting a diminished impact. For the Swiss Franc, on the other hand, the coefficient on the TED spread started to increase consistently from

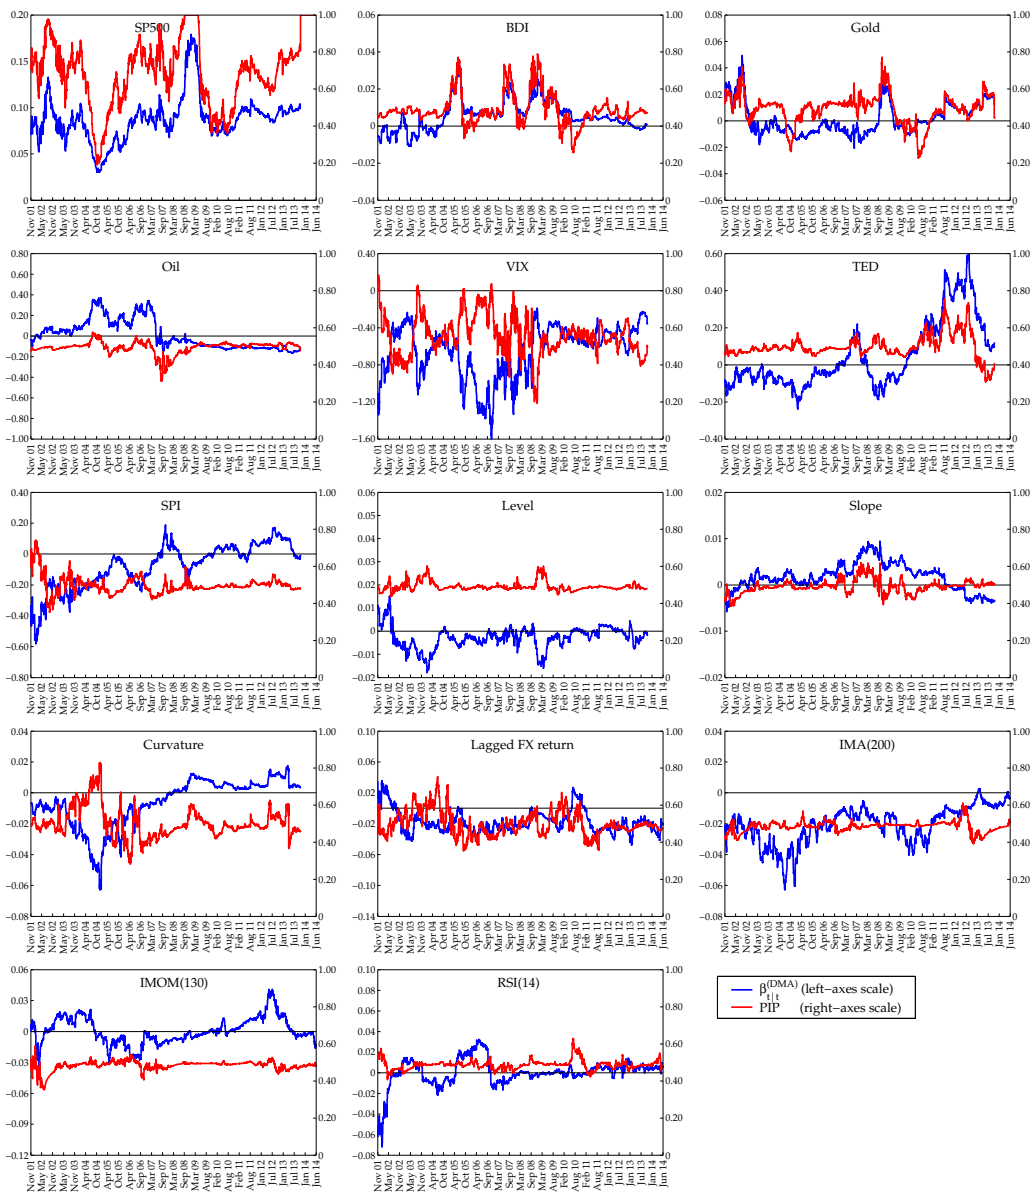


Fig. 8. Time series evolution of the (DMA) model averaged $\hat{\beta}_t$ estimates (blue solid line, left-axis scale) and the posterior inclusion probabilities (PIPs) (red solid line, right-axis scale) for the CHF series over the out-of-sample forecast evaluation period from November 23, 2001 to June 30, 2014. The plots show the parameter estimates of all the fundamentalist (plots 1 to 10) as well as the chartist predictor variables (plots 11 to 14), respectively. The fundamentalist based forecasts end on October 22, 2013 (due to the yield curve data). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

February 2010, then sharply from March 2011 until about September 2011 and did not start to decrease in magnitude until about August 2012. The VIX coefficient was largest (in absolute value) during the September 2006 to September 2007 period.

Third, the magnitude of the inclusion probabilities for a number of the predictor variables appears to be rather constant at values close to 50% for most of the out-of-sample evaluation period. This indicates that these variables are included in half of the best performing models that are constructed, meaning that the weighting scheme used for each of the $M = 2^k$ possible models is very close to equal weighting, i.e., $\pi_{it}^{(m)} = 1/M$. One needs to be careful here to not interpret this result as an indication of un-informativeness of the variables that are included in the predictor set. Even with an equally weighted combination approach, the real benefit of model averaging is that it implies a certain type of 'shrinkage' estimation (see the discussion on page 845 in Rapach et al., 2010 for a specific application to forecasting the US equity premium or Elliott et al., 2013 for a broader discussion in the context of complete subset regression).¹³ As with all shrinkage type estimators, the improvement in out-of-sample forecast performance is achieved by finding an optimal trade-off between (squared) bias and variance which define the MSFE. Thus, even predictor variables with posterior inclusion probabilities of 'only' around 50% contribute to an overall lowering of the MSFE and thus an improvement in the out-of-sample forecast performance.¹⁴

Fourth, the information content in the chartist predictor variables is rather heterogeneous over the 6 currencies. For instance, the RSI(14) indicator is sizeable for the Australian and Canadian Dollars, but less informative for the other 4 currencies, showing a mild increase in magnitude for the JPY and GBP series over the January 2012 to July 2012 period. The momentum indicator is important for the EUR and AUD series at the beginning of the out-of-sample period, while the moving average indicator increases in absolute importance during the September 2008 to January 2009 period for the Japanese Yen. The one period lagged exchange rate return has a sizeable negative coefficient from November 2003 to September 2008 for the AUD series and from February 2011 to about January 2013 for the Canadian Dollar, capturing a conditional feedback effect, i.e., a reversal given an appreciation (or depreciation) of the currency.

4.3.4. Forecast performance over sub-periods

From the time series plots of the cumSFE series shown in Fig. 1, it is clear that there are substantial differences in the forecast performance of the model combined forecasts over time. To understand how these visual differences translate into statistical differences, we re-examine the statistical forecast evaluation results over three separate sub-periods. These are as follows: (1) from November 23, 2001 to September 1, 2008, which we call the pre Lehman Brothers collapse period; (2) from September 2, 2008 to March 1, 2009, which we consider as the period shortly before and following the Lehman Brothers collapse; and (3) from March 2, 2009 to June 30, 2014 (October 22, 2013 for the Japanese Yen and the Swiss Franc), which is our post Lehman Brothers collapse period. These results are reported in Table 2. Table 2 is organized in 3 blocks, with each block corresponding to one sub-period. Under the heading MSFE/rel. in the first column of each block we show the MSFE of the random walk model's predictions in the 'Random Walk (RW)' row and the remaining 3 row entries are relative MSFEs of the chartist, fundamentalist and model combined forecasts, respectively, with the MSFE

¹³ Note here also that equally weighted forecast combinations are often found to outperform more complicated weighting schemes in practice (see Smith and Wallis, 2009 and Claeskens et al., 2015 for recent discussion of this), particularly in smaller sample sizes, due to estimation errors of the combination weight. In a somewhat different forecasting context, Koop and Korobilis (2013, p. 190) also obtain model probabilities that largely remain in the 0.3 and 0.5 interval.

¹⁴ We should also add here that the whole point of model averaging (or combination), as emphasised recently by Rapach et al. (2010), is that one wants to approximate the often complex data generating process for asset prices by a combination of models, each of which may be informative for prediction at some point in time, but not consistently so over the full data period that one is analysing. Moreover, as recently shown analytically by Huang and Lee (2010), even in the perfect setup of exogenous regressors and a stable data generating process (or forecasting environment), combination forecasts can outperform forecasts from the full regressor (or the 'kitchen-sink') model, in finite samples.

Table 2

Out-of-sample forecast evaluation over three sub-periods.

Model	MSFE/rel.	$R_{os}^2(\%)$	CW-stat.	p -Value	MSFE/rel.	$R_{os}^2(\%)$	CW-stat.	p -Value	MSFE/rel.	$R_{os}^2(\%)$	CW-stat.	p -Value
Sub-sample period	23.09.2001–01.09.2008 ($T_{os} = 1767$)				02.09.2008–01.03.2009 ($T_{os} = 130$)				02.03.2009–30.06.2014 ($T_{os} = 1390$)			
EUR												
Random walk (RW)	0.3171				1.3690				0.3779			
Chartist (C)	1.0048	-0.4761	-0.5839	0.7204	0.9976	0.2442	1.3573	0.0873	0.9999	0.0061	0.5710	0.2840
Fundamentalist (F)	0.9821	1.7908	5.4496	0.0000	0.9655	3.4533	1.9077	0.0282	1.0041	-0.4124	-0.2884	0.6135
Model combined (MC)	0.9807	1.9294	5.5682	0.0000	0.9621	3.7906	1.9967	0.0229	1.0126	-1.2587	-1.0379	0.8503
JPY												
Random walk (RW)	0.3391				1.6746				0.4199			
Chartist (C)	0.9995	0.0535	0.9375	0.1742	1.0105	-1.0458	-1.1412	0.8731	1.0010	-0.1009	0.0317	0.4874
Fundamentalist (F)	0.9943	0.5697	3.0784	0.0010	0.9795	2.0507	1.7346	0.0414	0.9983	0.1717	1.3064	0.0957
Model combined (MC)	0.9933	0.6684	3.5747	0.0002	0.9712	2.8814	1.9344	0.0265	0.9981	0.1873	1.5025	0.0665
GBP												
Random walk (RW)	0.2515				1.6475				0.2922			
Chartist (C)	1.0004	-0.0425	0.7270	0.2336	1.0078	-0.7781	-1.7781	0.9623	0.9993	0.0689	0.7992	0.2121
Fundamentalist (F)	0.9929	0.7060	3.3732	0.0004	0.9669	3.3109	2.3816	0.0086	1.0021	-0.2114	0.6713	0.2510
Model combined (MC)	0.9921	0.7920	3.6527	0.0001	0.9578	4.2240	2.5248	0.0058	1.0113	-1.1258	-0.5986	0.7253
CAD												
Random walk (RW)	0.2698				1.7956				0.3611			
Chartist (C)	1.0051	-0.5060	-1.3407	0.9100	0.9959	0.4128	1.8840	0.0298	1.0060	-0.6022	-2.0204	0.9783
Fundamentalist (F)	0.9912	0.8796	3.4587	0.0003	0.9953	0.4679	0.6198	0.2677	1.0044	-0.4376	-0.4502	0.6737
Model combined (MC)	0.9902	0.9775	3.7107	0.0001	0.9923	0.7706	0.8274	0.2040	1.0075	-0.7502	-0.5325	0.7028
CHF												
Random walk (RW)	0.4089				1.2511				0.5461			
Chartist (C)	1.0058	-0.5758	-1.3631	0.9136	0.9956	0.4375	0.9772	0.1642	1.0016	-0.1646	-0.0535	0.5213
Fundamentalist (F)	0.9831	1.6897	5.2492	0.0000	0.9203	7.9687	2.6115	0.0045	1.0082	-0.8244	-0.6014	0.7262
Model combined (MC)	0.9817	1.8284	5.5134	0.0000	0.9186	8.1359	2.6232	0.0044	1.0127	-1.2690	-0.8340	0.7978
AUD												
Random walk (RW)	0.4440				5.5462				0.6353			
Chartist (C)	1.0031	-0.3055	-0.2791	0.6099	0.9990	0.1006	0.5151	0.3033	0.9994	0.0635	0.9033	0.1832
Fundamentalist (F)	0.9849	1.5105	4.8523	0.0000	1.0089	-0.8854	-0.4063	0.6577	1.0006	-0.0620	0.0918	0.4634
Model combined (MC)	0.9848	1.5209	4.7983	0.0000	1.0114	-1.1389	-0.0055	0.5022	1.0081	-0.8080	-1.3792	0.9161

Notes: This table reports one-step ahead out-of-sample forecast evaluation results for the 6 exchange rates of interest over three different sub-periods. These are: (1) pre Lehman Brothers collapse from November 23, 2001 to September 1, 2008; (2) Lehman Brothers collapse from September 2, 2008 to March 1, 2009 and (3) post Lehman Brothers collapse from March 2, 2009 to June 30, 2014 (October 22, 2013 for CHF and JPY). In the column under the MSFE/rel. heading (in each sub-sample block) we show the MSFE of the random walk model in the row corresponding to Random walk (RW), followed by the relative MSFEs of the chartist, fundamentalist and model combined forecasts, respectively, with the MSFE of the random walk being the denominator. The terms R_{os}^2 , CW-statistic and p -value are as defined in Table 1.

of the random walk model being in the denominator, ie., the benchmark model. The terms $R_{os}^2(\%)$, CW-stat., and p -value are again the Campbell and Thompson (2008) R_{os}^2 , the Clark and West (2007) MSFE adjusted t -statistic and its corresponding one-sided p -value, as defined before.

The evaluation results over the three sub-periods in Table 2 confirm statistically the visual assessment from Fig. 1. First, consider the results over the pre Lehman Brothers collapse period. The combined forecasts from our empirical heterogeneous agent model strongly reject the null hypothesis of the forecasts not being different from those of a RW model. The p -values corresponding to the Clark and West (2007) MSFE adjusted test are well below 0.001 for all MC forecasts. For the CHF and EUR series, the CW-statistic is over 5, while for the remaining exchange rate series it is over 3. In terms of the magnitude of the improvement, the Campbell and Thompson (2008) out-of-sample R^2 is as high as 1.93% and 1.83% for the EUR and CHF series. The lowest R_{os}^2 values, which are obtained for the JPY and GBP series, are still a sizeable 0.67% and 0.79% at a daily sampling frequency. From the fundamentalist and chartist results it can be seen that the majority of the improvement in the model combined forecasts is driven largely by the fundamental predictor variables.

Looking over the results covering the Lehman Brothers collapse period shown in the second block in Table 2 one can see evidence of increased predictability (relative to the RW benchmark model) from the substantially larger out-of-sample R^2 values. These are as high as 8.14% for the model combined forecasts of the Swiss Franc series and sizeable 4.22%, 3.79% and 2.88% for the GBP, the EUR and the JPY. Moreover, these improvements are statistically significant at the 1% and 5% levels for the CHF and GBP, and EUR and JPY series, respectively. Note here that there are only 130 observations available to compute the CW-statistic, so that the results are (statistically) 'seemingly weaker' than for the longer pre-Lehman Brothers collapse period, ie., the CW-statistic < 3 , despite the much larger R_{os}^2 values. For the Canadian Dollar, the improvement is weak at only 0.77% and statistically insignificant. For the Australian Dollar, the R_{os}^2 is negative, highlighting the poor out-of-sample performance of the model combined forecasts relative to the RW model during this period.

Finally, from the third sub-period covering the post Lehman Brothers collapse shown in the last results block we can see that for all exchange rates, except for the JPY series, the performance of the model combined forecasts drops off substantially relative to the RW benchmark, producing negative R_{os}^2 values. These results also corroborate the visual assessment from the cumSFE series provided in Fig. 1. Overall, we can conclude from the sub-period analysis that, as conjectured from the cumSFE time series plots, there exists strong statistical evidence of exchange rate predictability before the Lehman Brothers collapse. Moreover, during the time of the Lehman Brothers collapse, predictability with respect to a benchmark RW model actually increased, resulting in much larger out-of-sample R^2 values.

4.3.5. Economic evaluation of predictability

How much 'economic value' do the improvements in the model combined forecasts produce? In our setting, we have 6 risky assets (and one risk-free one) that need to be aggregated into a portfolio of risky assets and hence requires portfolio weights. To evaluate the economic significance of our model combined forecasts, we implement a standard dynamic asset allocation methodology, as recently used in Della Corte et al. (2009), Della Corte et al. (2011) and Li et al. (2015). The dynamic asset allocation model is implemented as follows.¹⁵ Consider a US based investor which aims to build a currency portfolio consisting of six currencies. Investing in US government bonds is assumed to be riskless, so that the yield on a US government bond is used as the risk-free interest rate, which we denote by r_t^f . The expected return on a foreign bond is equal to the expected exchange rate change plus the interest earned on the foreign bond. We take the 3 months interest rates, ie., $y_t^{(3)}$, used in the construction of the level, slope and curvature factors as our short-term rate for both US and foreign bonds.

The dynamic asset allocation strategy that we implement and which maximizes the expected portfolio return $\hat{r}_{t+1|t}^p$ at each point in time is obtained by solving the following optimisation problem:

$$\max_{\omega_t} \hat{r}_{t+1|t}^p = \omega_t' \hat{r}_{t+1|t} + (1 - \omega_t') r_t^f \quad (17a)$$

¹⁵ We closely follow the description in Della Corte et al. (2011) and Li et al. (2015) in this section.

$$\text{subject to } \sigma_p^* = (\bar{\omega}' \Sigma_{t+1|t} \bar{\omega}_t)^{1/2}, \tag{17b}$$

where \mathbf{r}_{t+1} is a 6×1 vector of exchange rate returns, $\hat{\mathbf{r}}_{t+1|t} = \mathbb{E}_t(\mathbf{r}_{t+1})$ is its respective conditional forecast formed at time t , $\mathbf{1}$ is a 6×1 vector of ones, σ_p^* is the target conditional volatility of the portfolio returns and $\Sigma_{t+1|t} = \mathbb{E}_t[(\mathbf{r}_{t+1} - \hat{\mathbf{r}}_{t+1|t})(\mathbf{r}_{t+1} - \hat{\mathbf{r}}_{t+1|t})']$ is the 6×6 conditional variance–covariance matrix of \mathbf{r}_{t+1} . The solution to the maximum expected return problem in (17) is given by the following 6×1 vector of risky asset weights:

$$\bar{\omega}_t = \Sigma_{t+1|t}^{-1} (\hat{\mathbf{r}}_{t+1|t} - \mathbf{u}_t^f) \frac{\sigma_p^*}{\sqrt{\mathcal{K}_t}} \tag{18a}$$

$$\mathcal{K}_t = (\hat{\mathbf{r}}_{t+1|t} - \mathbf{u}_t^f)' \Sigma_{t+1|t}^{-1} (\hat{\mathbf{r}}_{t+1|t} - \mathbf{u}_t^f). \tag{18b}$$

The realised return on the investor's portfolio is then computed as:

$$\mathbf{r}_{t+1}^p = \bar{\omega}_t' (\mathbf{r}_{t+1} - \mathbf{u}_t^f) + r_t^f. \tag{19}$$

We evaluate the performance of our agent combined forecasts using the [Goetzmann et al. \(2007\)](#) performance measure, defined as

$$\mathcal{G}(r^p) = \frac{1}{(1-\gamma)} \ln \left\{ \frac{1}{T_{os}} \sum_{t=T_{is}}^{T_{fe}} \left(\frac{1 + \mathbf{r}_{t+1}^p}{1 + r_{t+1}^f} \right)^{(1-\gamma)} \right\}, \tag{20}$$

where γ is the parameter that captures the investor's relative risk aversion. The term $\mathcal{G}(r^p)$ is frequently interpreted as the certainty equivalent of the excess portfolio returns and does not require any assumptions regarding utility functions for a ranking of the portfolios to be constructed. We compare the performance of the agent combined forecasts to the RW benchmark by constructing the difference

$$\mathcal{P} = \mathcal{G}(r_{MC}^p) - \mathcal{G}(r_{RW}^p), \tag{21}$$

where $\mathcal{G}(r_{MC}^p)$ and $\mathcal{G}(r_{RW}^p)$ denote the [Goetzmann et al. \(2007\)](#) performance measure computed on portfolios constructed from the agent combined forecasts and the random walk model, respectively. The performance measure \mathcal{P} defined in (21) can be interpreted as the maximum performance fee that an investor is willing to pay to have access to the agent combined forecasts (relative to the RW benchmark). We follow the mainstream literature and report \mathcal{P} in annualised basis points (*bps*). We also report standard (annualised) Sharpe ratio's (*SRs*).

The economic evaluation results are reported in [Table 3](#). We only consider the out-of-sample period for which data for all 6 exchange rates are available, that is, from November 23, 2001 to October 22, 2013. We follow [Li et al. \(2015\)](#) and set the relative risk aversion parameter γ to 6 and the target volatility of the portfolio σ_p^* to 10%. We further set the covariance matrix $\Sigma_{t+1|t}$ in (17) and (18) at the unconditional covariance matrix (Σ) of exchange rate returns computed over the first 500 + 252 in-sample data points and do not update it. In line with the motivation given by [Li et al. \(2015\)](#), the reason for setting $\Sigma_{t+1|t} = \Sigma$ is that we want to isolate the effect of the conditional mean forecasts on the portfolio performance and do not want to confound it by allowing for a time varying covariance matrix. From the results in [Table 3](#) we can see that the model combined forecasts as well as the fundamentalist forecasts produce high economic value with (annualised) Sharpe ratios of 0.90 and 0.89, respectively. The performance fee of both the model combined and fundamentalists forecasts are over 400 annual basis points (relative to the RW model), suggesting that a risk averse investor (with $\gamma=6$) is willing to pay a fee of up to say 4% per year to have access to the model combined or fundamentalist forecasts. This finding is broadly in line with the magnitudes reported in [Li et al. \(2015\)](#), who use monthly data, nevertheless, relying on standard macro-economic variables that are not available in real time to produce the forecasts.

Table 3
Economic evaluation of the out-of-sample forecasts.

Investment strategy	$\mu_p(\%)$	$\sigma_p(\%)$	SR	$\mathcal{P}(bps)$
Random walk (RW)	8.0630	9.3230	0.5825	–
Chartist (C)	9.4325	11.6532	0.5622	53
Fundamentalist (F)	19.1128	16.6270	0.9007	422
Model combined (MC)	20.3030	18.2361	0.8925	467

Notes: This table reports the economic value of the various forecasting models that are considered. These are constructed over the out-of-sample period from November 23, 2001 to October 22, 2013 for all six currencies that we consider. We build maximum expected return portfolios, using a relative risk aversion parameter γ of 6 and a target portfolio volatility level of $\sigma_p^* = 10\%$. We use the first 500 + 252 observations from January 6, 1999 to November 22, 2001 to compute the variance covariance matrix of the exchange rate returns Σ needed to construct the portfolio weights in (18) and do not update Σ as we roll through the out-of-sample period. In columns two to five above, we show the annualised mean μ_p (in percent), volatility σ_p (in percent), Sharpe ratio SR and performance fee \mathcal{P} (in basis points, *bps*), respectively.

What is also interesting to see from the results in Table 3 is that the performance of chartist and random walk forecasts in terms of economic value are very similar, with the chartist's model producing somewhat higher mean returns than the RW, albeit at the cost of generating more volatility in the portfolio returns, thereby yielding a lower Sharpe ratio of 0.5622. The model combined forecasts also produce marginally higher portfolio volatility than a simple fundamentalist model portfolio, despite the higher mean return. Note here that the random walk portfolio is constructed by using the no-change forecast of the exchange rate plus the yield on the foreign bond as the predicted return in the weight construction in (18). That means that the random walk portfolio is based on weights from a carry trade strategy (see also page 4 in Li et al., 2015 for further intuition about this result). This last fact is important to highlight, as the gains in economic performance are not simply driven by a pure carry trade strategy.

In Fig. 9 we plot the cumulative wealth of the dynamic out-of-sample investment strategy for a portfolio with an initial wealth of 1 US Dollar to provide a visual impression of the portfolio performance from the various forecasting models that we consider. The cumulative portfolio wealth constructed from the model combined forecasts (blue line) and also the fundamentalist forecasts only (red line) show a similar pattern to what we observed from the cumulative sum of squared forecast errors in Fig. 1. For instance, there is an overall upward trend in cumulative wealth from the beginning of the out-of-sample portfolio construction period until September 2008, i.e., up to the time of the Lehman Brothers collapse. From September 2008 to about February 2009, there is a substantial increase in cumulative wealth, and from the end of February 2009, cumulative wealth remained fairly constant until the end of the sample in October 2013. The pure chartist and random walk (carry trade portfolios) perform fairly similarly in the sense that at the portfolio level, the chartist weights are driven largely by the interest differential as opposed to the superior predictions from the technical trading rules. Nevertheless, there do appear to be instances where the technical trading rules perform better than a random walk based portfolio (see for instance, the period from May 2010 to August 2011).

4.4. Forecast performance at longer horizons

Given the overall positive results at the one day ahead horizon, we now turn to assess the forecasting performance of our proposed fundamentalist and chartist model combined predictions at horizons of 2 up to 5 days ahead. Here we simply focus on a statistical evaluation to conserve space and avoid repetition. To construct multiple-step ahead out-of-sample forecasts from the agent based model combination, we implement the so-called 'direct' forecasting approach. That is, we re-formulate the relation in (9) for the general h -step ahead relation as:




Fig. 9. Cumulative portfolio wealth for the different forecasting models. Initial wealth is set to 1 US Dollar. The cumulative wealth is constructed from the dynamic investment strategy using the out-of-sample forecasts to construct the portfolio weights. The out-of-sample forecasts are constructed from the model combined (blue line), fundamentalist (red line), chartist (green line) and random walk model (black dashed line), respectively, over the out-of-sample period from November 23, 2001 to October 22, 2013. The horizontal line at 1 marks the initial wealth of 1 US Dollar. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$\hat{r}_{t+h|t}^{MC} = (1 - \hat{\omega}_{t,h}) \hat{r}_{t+h|t}^F + \hat{\omega}_{t,h} \hat{r}_{t+h|t}^C \tag{22}$$

where $\hat{r}_{t+h|t}^F$ and $\hat{r}_{t+h|t}^C$ are respectively, the individual h -step ahead (h -period holding return) fundamentalist and chartist (direct) forecasts. The h -step ahead weighting function $\hat{\omega}_{t,h}$ is obtained, analogous to (11), from the regression:

$$(r_{t-h:t} - \hat{r}_{t|t-h}^F) = \omega_{t,h} (\hat{r}_{t|t-h}^C - \hat{r}_{t|t-h}^F) + v_t \tag{23}$$

with $r_{t-h:t} = 100(P_t/P_{t-h} - 1)$ being the h -period holding return of the currency of interest.¹⁶

4.4.1. Multiple-step ahead forecast evaluation

We use the same calibration for the λ , α and κ parameters that were used in the one-step ahead prediction setting to implement the Kalman Filter recursions. One could again try to optimise over these parameters, but due to the computational burden, we do not consider this here. The results for forecast horizons $h = 2, 3, 4$, and 5 days ahead are reported in Table 4. To avoid clutter in the table, we only report results for the model combined forecasts for each exchange rate series at the four different forecast horizons.¹⁷ Also, since h -step ahead forecast errors will be $MA(h - 1)$ processes in general

¹⁶ How the agent specific multiple period ahead out-of-sample forecasts are computed is described in more detail in the Appendix.

¹⁷ At horizons 2 to 5, we also find that for the majority of exchange rates, the combined forecasts improve on the fundamental ones. The exception is again the JPY series, at all forecast horizons and also the CAD series for forecast horizons above 2. Further details, if needed, are available upon request from the authors.

Table 4
Multiple-steps-ahead out-of-sample forecast evaluation results.

Forecast horizon	MSFE _(RW)	MSFE _(MC)	Relative-MSFE	R _{os} ² (%)	CW-statistic	p-Value
EUR						
h = 2	0.7844	0.7817	0.9965	0.3519	2.3546	0.0093
h = 3	1.1610	1.1574	0.9969	0.3105	1.8488	0.0322
h = 4	1.5226	1.5271	1.0030	-0.2961	1.2312	0.1091
h = 5	1.8956	1.9077	1.0064	-0.6389	0.8116	0.2085
JPY						
h = 2	0.8227	0.8206	0.9975	0.2510	2.3214	0.0101
h = 3	1.2069	1.2091	1.0019	-0.1869	0.9743	0.1650
h = 4	1.5956	1.6071	1.0072	-0.7208	0.7506	0.2265
h = 5	1.9627	1.9731	1.0053	-0.5294	0.9648	0.1673
GBP						
h = 2	0.6620	0.6611	0.9986	0.1355	1.8221	0.0342
h = 3	0.9923	0.9945	1.0022	-0.2206	1.0510	0.1466
h = 4	1.3243	1.3236	0.9995	0.0547	0.9869	0.1618
h = 5	1.6493	1.6537	1.0026	-0.2648	0.5734	0.2832
AUD						
h = 2	1.3889	1.3943	1.0039	-0.3938	0.6346	0.2628
h = 3	2.0339	2.0324	0.9992	0.0756	1.4167	0.0783
h = 4	2.6399	2.6402	1.0001	-0.0110	1.0498	0.1469
h = 5	3.3072	3.3197	1.0038	-0.3771	0.9554	0.1697
CAD						
h = 2	0.6689	0.6672	0.9974	0.2639	1.3240	0.0928
h = 3	0.9955	0.9927	0.9972	0.2823	1.7542	0.0397
h = 4	1.3267	1.3277	1.0008	-0.0778	0.8242	0.2049
h = 5	1.6831	1.6846	1.0009	-0.0890	0.8105	0.2088
CHF						
h = 2	0.9598	0.9581	0.9982	0.1828	2.2541	0.0121
h = 3	1.4196	1.4194	0.9999	0.0126	1.4049	0.0800
h = 4	1.8704	1.8732	1.0015	-0.1475	1.1499	0.1251
h = 5	2.3059	2.3161	1.0044	-0.4447	0.8619	0.1944

Notes: This table reports multiple-step ahead out-of-sample forecast evaluation results for the 6 exchange rates of interest. In column 1, the various forecast horizons that are considered in the evaluation are listed. Columns 2, 3 and 4 show the mean squared forecast errors (MSFEs) of the random walk (RW) benchmark, the MSFEs of the model combined (MC) predictions, and the Relative-MSFEs defined as $MSFE_{(MC)}/MSFE_{(RW)}$. In column 5 we report the Campbell and Thompson (2008) out-of-sample R^2 (denoted by R_{os}^2) which is computed as $1 - MSFE_{(MC)}/MSFE_{(RW)}$. The Clark and West (2007) CW-statistic and its corresponding p-value are given in columns 6 and 7. The out-of-sample evaluation period is from November 26, 2001 to June 30, 2014 for all currencies, except for the JPY and CHF series, where the sample ends again in October 22, 2013. We use a HAC robust variance in the computation of the CW-statistic, employing a Quadratic Spectral Kernel on the pre-whitened cw_{t+1} series, using an ARMA(1,1) model as the approximating model and choose a data driven AR(1) bandwidth to compute the HAC of the pre-whitened series.

(ie., will be moving averages and therefore autocorrelated of order $h - 1$) which affects the CW-statistic, we use a HAC robust variance for $\text{Var}(\bar{c}w)$ in (14). More specifically, we follow the recommendation of Andrews and Monahan (1992) and employ a data driven bandwidth using a Quadratic Spectral (QS) Kernel with a 'pre-whitening' step, where we choose the (optimal) bandwidth parameter with an AR(1) as the approximating model (see equation 3.5 in Andrews and Monahan, 1992).¹⁸ The last four columns in Table 4 are the same as the last four columns in Table 1. Columns 2

¹⁸ That is, to pre-whiten the cw_{t+1} series, we first fit an ARMA(1, 1) to cw_{t+1} and then use the QS Kernel with the bandwidth parameter set to $1.3221(\hat{\alpha}(2)T_{os})^{1/5}$, where $\hat{\alpha}(2) = 4\hat{\rho}^2/(1-\hat{\rho})^4$ and $\hat{\rho}$ is the AR(1) parameter estimate obtained from an AR(1) regression of the (pre-whitened) residual series obtained from the ARMA(1, 1) model fitted to cw_{t+1} . To obtain the HAC variance, we then 're-colour' again with the ratio of the square of the ARMA lag polynomials (see Andrews and Monahan, 1992 for more details on the exact computations).

and 3 in Table 4 show respectively the MSFEs of the RW benchmark and the MC predictions. The first column lists the various forecast horizons.

From the results reported in Table 4 it is evident that the forecast performance of the MC predictions – relative to the random walk forecast – remains intact for forecast horizons up to 2 days ahead for all currencies (except for the Australian Dollar) generating out-of-sample R^2 values of 0.35, 0.25, 0.14, 0.26 and 0.18 percent for the EUR, JPY, GBP, CAD and CHF exchange rate series, respectively. Moreover, these improvements are statistically significant at the 5% level for the EUR, JPY, GBP, and CHF, and at the 10% level for the CAD. What is interesting to see is that for a holding period return of even 3 days, improvements in the R_{os}^2 are still realised for the EUR ($R_{os}^2 = 0.31\%$), the CAD ($R_{os}^2 = 0.28\%$) and also for the AUD series, albeit with a lower R_{os}^2 of 0.08, and are overall significant at the 10% level. For the CAD and the AUD series, forecasting the three day holding return $r_{t:t+3}$ yields in fact higher R_{os}^2 values than forecasts two days ahead. At forecast horizons of 4 and 5 days ahead, the MC predictions start to produce consistently worse forecasts than the benchmark random walk model across all 6 exchange rates, as the forecast horizon increases.

We again show plots of the cumSFE of the model combined forecasts, for $h = 2, \dots, 5$ in Fig. 10 to visualise how the performance of the forecasts evolves over the out-of-sample evaluation period relative to the RW benchmark. In order to facilitate the comparison across the various forecast horizons, we plot all cumSFE for each exchange rate and considered forecast horizon in one subfigure in Fig. 10 and offset the various horizons so that they can be plotted in the same graph. From the time series plots of the cumSFE series in Fig. 10, a number of interesting features stand out. First, the increase in the cumSFE around (and following) the time of the Lehman Brothers collapse remains visible for forecast horizons of up to 5 days ahead for the EUR, GBP, CAD series, and also, but to a lesser extent, for the CHF series. For the Australian Dollar, a substantial worsening in the predictive ability with respect to the RW benchmark can be seen, most evidently for forecast horizons 5, 4 and 2 days ahead. What is perhaps somewhat surprising to see from the cumSFE plot for the Australian Dollar is the improved forecast result at the 3 days ahead horizon. This is evident from the (red) cumSFE corresponding to $h = 3$ being consistently above the 0 line in Fig. 10. The Lehman Brothers collapse seems to have had a rather positive impact on the predictive performance for 3 day ahead forecasts, and is therefore more inline with the effects experienced by the other 5 currencies.

Second, the improved forecast performance at the 3 steps-ahead horizon which were found from the statistical evaluation results in Table 4 for the EUR and CAD are largely driven by the strong performance of the MC forecasts around the time of the Lehman Brothers collapse. This can be seen from the persistent upward jump in the cumSFE around the September 2008 period. What is interesting to highlight here is that the forecast performance of the MC predictions remained fairly stable after the Lehman Brothers collapse, which was not the case for the Canadian Dollar at the one-step ahead horizon, while for the EUR the cumSFE series decreased somewhat, indicative of a mild worsening in forecast performance with respect to the RW benchmark. Third, for the Swiss Franc, the MC predictions are consistently superior to RW forecasts over the period from September 2008 to September 2011 for all 5 forecast horizons that we consider. Moreover, a noticeable build-up in the predictive improvement seems to occur from approximately November 2010 until about August 2011.

5. Conclusion

We build an empirical heterogeneous agent model consisting of fundamentalists and chartist agents to forecast 6 of the most frequently traded currencies using daily data over the sample period from January 1999 to June 2014. More specifically, we use a time varying model combination approach to optimally average the forecasts from individual fundamentalist and chartist agents, where individual fundamentalist and chartist predictions are constructed using the recently proposed flexible DMA framework. We use daily financial data as proxy variables for the macroeconomic predictors used by fundamentalist agents. These fundamental predictors contain level, slope and curvature yield curve factors, equity indices as well as data related to global trade activity and risk aversion. To model the behaviour of chartist agents, we construct technical predictor variables consisting of moving averages and momentum to capture the well known trend following behaviour of chartists and also relative

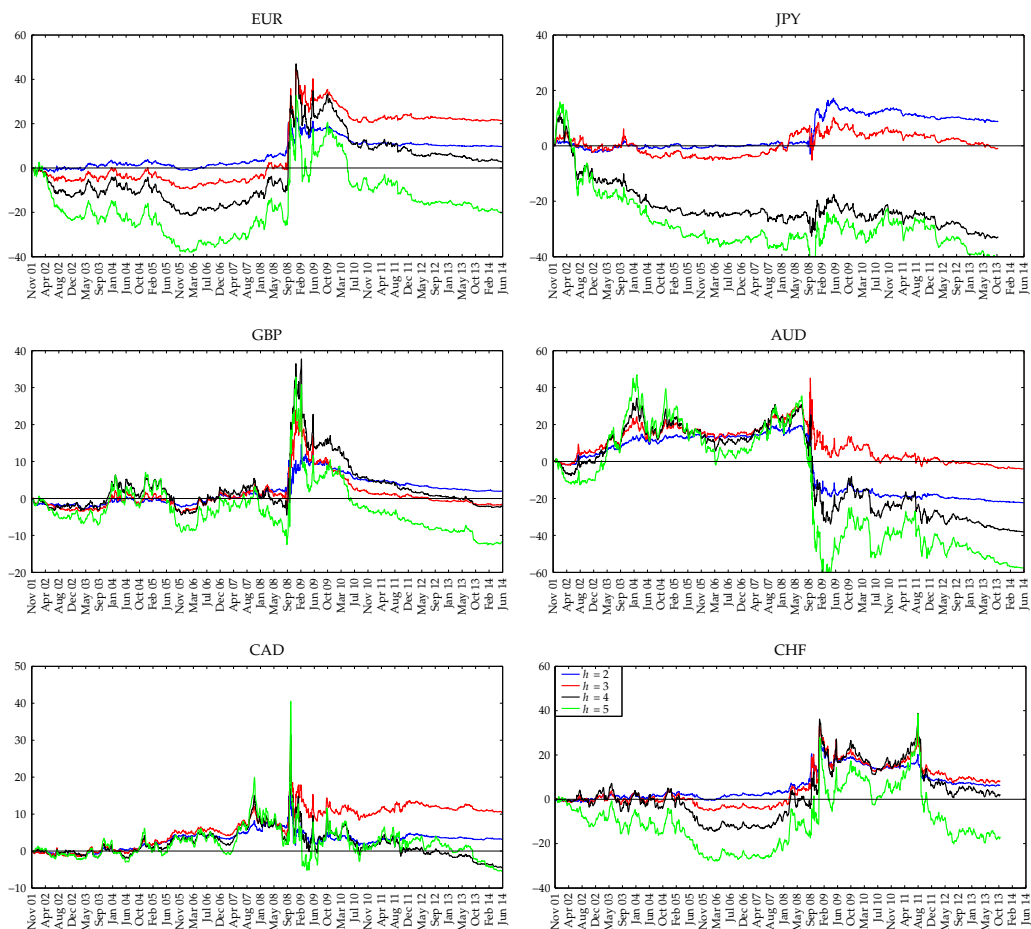


Fig. 10. Time series evolution of the h -step ahead cumulative squared forecast errors (cumSFE_t) of the model combined forecasts (relative to the RW model) for forecast horizons $h = 2, 3, 4$ and 5 over the out-of-sample period from November 26, 2001 to June 30, 2014 (October 22, 2013 for the CHF and JPY series). Note that we plot all series up to June 30, 2014 so that the dates on the x -axis can be compared easily. Also, we offset the forecasts for $h = 3, 4$ and 5 with 1 to three dimensional NAN vector entries in the forecasts so that the dimensions of the forecast vectors are the same and can be compared on the same time scale. (For interpretation of the references to colour in this figure, the reader is referred to the web version of this article.)

strength indices to capture overbought or oversold conditions in the foreign exchange market and reversal trading strategies.

We show that out-of-sample forecasts from our empirical heterogeneous agent model significantly outperform the forecasts from a random walk benchmark model for all 6 currencies that are considered. Out-of-sample R^2 values can be as high as 1.41%, 1.07%, 0.99%, and 0.74% for the Franc, the Euro, the Pound and the Yen series, and are somewhat lower for the Australian and Canadian Dollars at 0.29% and 0.24%. Statistical tests show that these forecast gains are significant at the 10% level for the Australian and Canadian Dollars, and at the 1% level for the remaining 4 currencies. Moreover, using a dynamic asset allocation framework we show further that our model combined forecasts also generate economic value. Annualized Sharpe ratios are as high as 0.89, yielding performance fees of over 460 (annualized) basis points, relative to a random walk benchmark model. Forecast gains remain statistically significant for forecasts up to three days ahead for some currencies.

We show further that there is substantial instability in the predictive performance of our model. The majority of the forecast gains are realised before and during the Lehman Brothers collapse period, that is, from the beginning of the out-of-sample period until February 2009. This is visible from the cumulative squared forecast errors as well as the cumulative wealth of a portfolio formed with a dynamic asset allocation strategy. These results are also confirmed with statistical tests on sub-samples. Although the modelling approach that we adopt in this paper is extremely flexible, as it can accommodate time-varying parameters as well as time varying predictor variables, the model is not overly suitable for the post February 2009 period, after which the predictive performance worsens.

Possible avenues for future research could be to expand the set of exchange rates to see how well the model performs at an even wider cross-section. Also, it would be interesting to see if the same instabilities are evident for a wider set of currencies, thereby suggesting that there is a common risk factor that lead to this breakdown in predictive performance. It would also be interesting to see if a re-calibration of the forgetting factors (α and λ) that control the discounting of the data in the model and thereby the number of observations that are used for estimation improve the forecast performance in the post February 2009 period. The period around the Lehman Brothers collapse could alternatively be viewed as an 'outlier' time period, and one could work with different weighting functions to discount or 'robustify' this effect on the predictive performance of the model.

Appendix

This appendix provides additional details on the construction of the fundamental and technical indicators, the DMA modelling framework as well as further results with regard to the calibration of the α and λ to supplement the findings in the main part of the paper. The content of this appendix are as follows:

Contents of Appendix: A-1

- A.1. Description of data and variables A-2
 - A.1.1. Summary statistics of exchange rates and fundamental predictors A-2
- A.2. Construction fundamental variables and technical indicators A-6
 - A.2.1. Information in the yield curve A-6
 - A.2.2. Information in stock prices A-7
 - A.2.3. Risk aversion measures and global trade activity A-8
 - A.2.4. Moving Average Rules A-9
 - A.2.5. Momentum Indicators A-10
 - A.2.6. Relative Strength Index A-10
- A.3. Details of Dynamic Model Averaging (DMA) framework A-11
- A.4. Calibration of λ and α A-15
- A.5. Computing agent specific multiple period ahead out-of-sample forecasts A-16

A.1. Description of data and variables

Here we briefly describe some of the basic features of the exchange rate data and the two different sets of predictor variables that we use in our forecast evaluation. Note that the discussion here is not meant to be exhaustive, but rather complementary to the tables and figures that we provide to summarise this information.

We use simple returns in all our return calculations, rather than log-returns computed from differences in the log prices. That is, returns are computed as $r_t = 100(P_t/P_{t-1} - 1)$, where P_t is the time t price of the asset of interest. Our motivation for using (simple) returns is trivial. Due to the financial crisis period, there were at times substantial daily price variations, particularly for the BDI, gold prices, equity prices and most evidently for oil prices. For instance, on the 24th of September 2001, the log-return for (WTI) oil was -16.55 , while the simple return was -15.25 , a difference of 1.3 percentage

points. Similarly, on the 29th of December 2008, the log-return and simple return were 21.28 and 23.71, respectively, a difference of around 2.4 percentage points. To be able to capture the 'true' daily variations an investor was exposed to, we prefer to use the simple return (just return henceforth) construction. We should stress here also that our predictability results are not affected by the specification of the return process and hold equally well when log-returns are used.

A.1.1. Summary statistics of exchange rates and fundamental predictors

In Fig. A1 we show plots of the 6 different currencies that are used in the forecast evaluation exercise over the full sample period from January 4, 1999 to June 30, 2014. The left column in Fig. A1 shows the (raw) US Dollar price of one foreign currency unit, and the right column shows the daily returns of the series. As a reminder, an upward movement in these 6 series indicates that the respective currency has appreciated against the US Dollar (the US Dollar price of the foreign currency has risen), while a downward movement suggests a depreciation. There are a number of interesting visual features that are evident from the plots in Fig. A1. First, notice how all 6 exchange rates show a general upward trend, suggesting that the series has appreciated over the last 15 years against the US Dollar. This trend is much weaker for the British Pound series from mid September 2008 onwards. Second, the Lehman Brothers collapse in September 2008 had a rather profound effect on the British Pound, the Australian Dollar, the Canadian Dollar and the Euro, resulting in depreciations of approximately 22%, 18%, 17%, and 10%, respectively, from September 1, 2008 to March 1, 2009.¹⁹ These four currencies thus behaved inline with what would be expected from an 'investment currency', where high levels of risk aversion lead to sell-offs in such a particular asset class.

Over the same time frame, the Japanese Yen appreciated by nearly 11%, while the Swiss Franc remained rather stable, depreciating only marginally by 2%. It is interesting to see here that the Japanese Yen behaved in accordance with its widely perceived 'safe haven' status, that is, providing financial refuge during times of high risk aversion, while the Swiss Franc, also known as a 'safe haven' currency, was largely unaffected. The Swiss Franc's 'safe haven' status did not come to bear any significant importance until the first set of problems began which eventually lead to the European sovereign bond crisis. The strongest Swiss Franc appreciation was realised over the July 2010 to August 9, 2011 period, where the currency surged over 41% against the US Dollar.²⁰ It is interesting to point out here that, although the European sovereign bond crisis appears to have been one of the key drivers of the 'safe haven' effect of the Swiss Franc, the strong appreciation in the Franc was not matched by an equal depreciation in the Euro against the US Dollar. The Euro, in fact, appreciated by nearly 17% over the same time period, highlighting that there must have been other additional factors that contributed to the strong appreciation of the Swiss Franc. The fact that the Euro did not depreciate against the US Dollar as a result of the European sovereign bond crisis could partly be due to the second round of the quantitative easing programme of the Federal Reserve being implemented in the US.

In the right column of Fig. A1, the time series evolution of the return series is plotted. As was the case in the levels plot of the 6 series, the homogenous response to the Lehman Brothers collapse in September 2008 is also clearly visible in all 6 return series of the currencies, with the response of the Swiss Franc, nevertheless, being somewhat weaker than for the other currencies. The Australian Dollar is the most volatile currency, with daily returns in the AUD swinging between -8 and 8 percent throughout most of October 2008. The intervention by the Swiss national bank via the imposition of the cap on the CHF/EUR rate at 1.20 is the most outstanding event impacting on the 6 return series after the Lehman Brothers collapse.

In Table A1 we show summary statistics of the 6 exchange rate returns to be forecasted (top 6 rows), as well as summary statistics of the fundamental predictor variables that are used by fundamentalist

¹⁹ For all 4 currencies, the depreciation relative to the US Dollar already started somewhat earlier; nevertheless, it is clear from the plots that from September 2008 the drops in the currencies amplified substantially.

²⁰ This surge may have been triggered by a number of events, starting with Greece needing "official" financial assistance in May 2010, followed by Ireland's bailout in November 2010 and with Portugal following in May 2011 (see Lane, 2012, p. 56). Note here also that the Swiss national bank imposed the cap on the CHF/EUR rate at 1.20 on September 6, 2011, but speculation about the implementation of the cap had already been circulating for weeks beforehand, so that the peak in fact occurred on August 9.




Fig. A1. Time series plots of the six exchange rates of interest (full sample period from January 4, 1999 to June 30, 2014). Left Panel shows the (raw) level series. Right Panel shows the return series used in the forecast evaluation.

Table A1
Summary statistics of all exchange rate returns and fundamental predictor variables.

Variable name	Description and data source	Mean	Median	Std. Dev.	Skew	Kurt	Min	Max	ACF(1)	
$r_{t,EUR}^i$	Exchange rate returns: $r_t^i = 100(P_t^i/P_{t-1}^i - 1)$, $\forall i$ {EUR, JPY, GBP, AUD, CAD, CHF}. All exchange rate data are taken from Bloomberg, where Prices P_t^i are PX_Last entries for the i th currency, using the New York 5:00 pm snap time stamp.	0.0056	0.0130	0.6359	0.0683	4.3251	-2.4870	3.5445	-0.0158	
$r_{t,JPY}^i$		0.0046	0.0082	0.6560	0.0316	6.4730	-5.3555	3.5659	-0.0411	
$r_{t,GBP}^i$		0.0023	0.0143	0.5591	-0.2377	5.2895	-3.4123	2.9633	0.0272	
$r_{t,AUD}^i$		0.0138	0.0507	0.8355	-0.2062	12.6416	-7.0500	8.6215	-0.0431	
$r_{t,CAD}^i$		0.0105	0.0129	0.5667	-0.1146	6.3163	-3.2017	4.0791	-0.0461	
$r_{t,CHF}^i$		0.0131	0.0088	0.6932	-0.3031	11.3561	-8.6881	4.8037	-0.0380	
$\Delta(L_t^{US} - L_t^{EU})$	Differences in Level, Slope and Curvature factors, relative to the US. Factors are constructed as described in (A.1). US yield curve data are taken from the Gürkaynak et al.'s (2007) database. For the Euro area, yields are combined from Bundesbank and ECB yield curve data, using the Nelson–Siegel–Svensson parameter estimates. Yield curve data for CHF and JPY are constructed using the Nelson–Siegel–Svensson parameter estimates from Malkhozov et al. (2014) . All yield curve data (except for CHF and JPY) were taken from the respective central bank websites at: Bundesbank: http://www.bundesbank.de/Navigation/EN/Statistics/Time_series_databases/Macro_economic_time_series/its_list_node.html?listId=www_s140_it03c ECB: https://www.ecb.europa.eu/stats/money/yc/data/fmd/download/yc_historical.zip RBA: http://www.rba.gov.au/statistics/tables/xls/f17hist.xls BoE: http://www.bankofengland.co.uk/statistics/Documents/yieldcurve/uknom05_mdaily.xls BoC: http://www.bankofcanada.ca/rates/interest-rates/bond-yield-curves/ .	-0.0002	0.0001	0.0470	-0.0145	9.7703	-0.4113	0.3947	-0.2300	
$\Delta(S_t^{US} - S_t^{EU})$		-0.0004	0.0000	0.0784	0.4713	10.6002	-0.3974	0.9412	-0.1841	
$\Delta(C_t^{US} - C_t^{EU})$		-0.0002	0.0000	0.0859	-0.2833	10.1404	-0.9713	0.6268	-0.1508	
$\Delta(L_t^{US} - L_t^{JP})$		-0.0007	-0.0003	0.0483	-0.1844	9.3525	-0.4275	0.3579	-0.1102	
$\Delta(S_t^{US} - S_t^{JP})$		-0.0009	0.0005	0.0876	-0.0429	8.6251	-0.7107	0.6840	-0.1101	
$\Delta(C_t^{US} - C_t^{JP})$		-0.0006	-0.0012	0.0891	0.0801	13.8899	-0.9807	0.8900	-0.1089	
$\Delta(L_t^{US} - L_t^{GB})$		-0.0001	0.0000	0.0440	-0.4554	10.1648	-0.4537	0.2788	-0.1825	
$\Delta(S_t^{US} - S_t^{GB})$		0.0003	0.0010	0.0659	0.0280	7.9994	-0.4737	0.4746	-0.0708	
$\Delta(C_t^{US} - C_t^{GB})$		0.0000	0.0000	0.0716	-0.0021	6.6027	-0.4346	0.4945	-0.0302	
$\Delta(L_t^{US} - L_t^{AU})$		-0.0005	0.0000	0.0605	0.1667	7.5503	-0.4581	0.4690	-0.3141	
$\Delta(S_t^{US} - S_t^{AU})$		-0.0003	0.0001	0.0849	0.1280	6.5926	-0.4861	0.7147	-0.2466	
$\Delta(C_t^{US} - C_t^{AU})$		-0.0002	0.0000	0.0871	0.0158	5.5323	-0.4674	0.4797	-0.0855	
$\Delta(L_t^{US} - L_t^{CA})$		-0.0001	0.0000	0.0337	-0.4586	13.2736	-0.4077	0.2628	-0.1100	
$\Delta(S_t^{US} - S_t^{CA})$		-0.0002	0.0000	0.0699	-0.1771	12.4698	-0.8162	0.5590	-0.1052	
$\Delta(C_t^{US} - C_t^{CA})$	-0.0001	0.0000	0.0961	-0.4087	86.6554	-1.9131	1.8418	-0.1531		
$\Delta(L_t^{US} - L_t^{HK})$	-0.0003	0.0001	0.0436	-0.3609	9.9029	-0.4370	0.3168	-0.2063		
$\Delta(S_t^{US} - S_t^{CH})$	-0.0003	0.0012	0.0694	-0.1668	7.4639	-0.5949	0.4574	-0.0893		
$\Delta(C_t^{US} - C_t^{CH})$	-0.0003	-0.0005	0.0718	0.1339	8.0946	-0.4149	0.5983	-0.0858		
$r_{t,SP500}^i$	Equity index returns: $r_t^i = 100(P_t^i/P_{t-1}^i - 1)$, $\forall j = \{DAX30, Nikkei225, FTSE100, AllOrds, SPTSX, SPI\}$. All equity price data are taken from Bloomberg. Prices P_t^i are PX_Last entries for the j th equity price.	0.0196	0.0224	1.2649	0.0155	11.1647	-9.0350	11.5800	-0.0818	
$r_{t,DAX30}^i$		0.0273	0.0486	1.5332	0.1312	7.7715	-8.4923	11.4020	-0.0167	
$r_{t,Nikkei225}^i$		0.0143	0.0000	1.5025	-0.2268	9.5657	-11.4064	14.1503	-0.0225	
$r_{t,FTSE100}^i$		0.0109	0.0000	1.2230	-0.0011	9.2853	-8.8493	9.8388	-0.0423	
$r_{t>AllOrds}^i$		0.0205	0.0247	0.9574	-0.4805	9.2634	-8.1980	5.5064	-0.0152	
$r_{t,SPTSX}^i$		0.0273	0.0409	1.1395	-0.4777	11.9827	-9.3242	9.8233	-0.0130	
$r_{t,SPI}^i$		0.0211	0.0259	1.1206	0.0059	9.5866	-7.2471	10.5788	0.0339	
ΔVIX_t		Difference/change in the VIX and TED spread. VIX data are taken from Bloomberg. TED spread (TEDRATE) is obtained from the St. Louis Fed FRED2 database http://research.stlouisfed.org/fred2/ . Returns on Gold, the BDI and Oil prices computed as $r_t^i = 100(P_t^i/P_{t-1}^i - 1)$, $\forall i = \{\text{Gold, BDI, Oil}\}$.	-0.0036	-0.0400	1.6668	0.5539	21.8168	-17.3600	16.5400	-0.1226
ΔTED_t			-0.0001	0.0000	0.0563	0.7999	68.3461	-0.8000	0.9900	0.1340
$r_{t,Gold}^i$			0.0445	0.0448	1.1507	0.0226	10.0596	-9.0737	10.7883	-0.0099
$r_{t,BDI}^i$	0.0187		0.0000	1.8285	0.0915	9.6616	-11.3716	14.6341	0.7992	
$r_{t,Oil}^i$	0.0810	0.0307	2.3666	0.1686	9.0281	-15.2542	23.7094	-0.0172		

Notes: This table reports standard summary statistics of the 6 exchange rate returns used in the forecast evaluation. These are shown in the first 6 rows at the top of the table. This table shows also summary statistics of all the fundamental predictor variables that are used in the forecasting model. These are shown below the main horizontal line in the middle of the table, starting from row 7, and are grouped according to the 3 broad groups that are defined in (A.2), (A.3) and (A.4). Dashed horizontal lines are used to separate the three broad groups of fundamental predictors. The abbreviation ACF(1) denotes the first order autocorrelation of the series. These summary statistics cover the period from from January 4, 1999 to June 30, 2014 for all exchange rate and predictor variables, except for the yield curve factors of Japan and Switzerland, which end already in October 22, 2013.

agents over the full sample period from January 4, 1999 to June 30, 2014.²¹ The fundamental predictor variables are arranged coherently in three separate blocks to match the description of the variables in the text. Looking over the summary statistics of the 6 exchange rates, a number of the stylised facts which were visible in Fig. A1 are also evident in the summary statistics. First, the mean return on the 6 exchange rates series is positive, confirming the positive trend in the level series and the overall depreciation of the US Dollar against these 6 currencies. Second, the Australian Dollar is the most volatile series, with a (daily) standard deviation of 0.84 percent, followed by the Swiss Franc with a standard deviation of 0.69 percent. These two currencies have further the 'heaviest tails', as is evident from their kurtosis statistics being well above 10.²² What is interesting to point out from the summary statistics of the exchange rate returns are the rather sizeable negative first order autocorrelations (henceforth ACF(1)) for the returns on the Yen, Canadian and Australian Dollars, as well as the Swiss Franc, as shown in the last column of Table A1.

Looking over the summary statistics of the fundamental predictor variables that are reported below the solid line separating the exchange return series from the predictors in Table A1, it is evident that the most volatile predictors are oil returns, BDI returns, as well as DAX30 and Nickkei225 equity returns. All equity returns as well as returns from investing in gold, oil and the BDI had positive means, raging from values as low as 0.01% for the return on the FTSE100 up to 0.08% for the return on oil. The return from investing in oil was not only the most volatile series, but it also had the single largest positive daily return of nearly 24%. From the first order ACFs in the last column of Table A1 it can be seen that all but one equity return series are negatively correlated, with the SP500 returns showing the strongest negative autocorrelation of -0.08. The single equity series with a positive ACF is the return on the Swiss SPI. By far the strongest first order autocorrelated predictor variable is the return on the BDI with an ACF(1) of 0.80.²³

A.2. Construction fundamental variables and technical indicators

This section elaborates on the construction of the fundamental and technical indicators that are used as predictor variables in the paper.

A.2.1. Information in the yield curve

The use of yield curve data is motivated by the findings in Bekaert and Hodrick (1992) and Clarida et al. (2003), who show that the information content in the yield curve is valuable for exchange rate forecasting.²⁴ Moreover, in the context of macro-finance models, Diebold et al. (2006) and Rudebusch and Wu (2008) have documented that the empirically derived level and slope factors are 'strongly' correlated with inflation and economic activity.²⁵

To capture the information in the yield curve, we construct level, slope and curvature factors, denoted by \mathcal{L}_t , \mathcal{S}_t and \mathcal{C}_t , using daily data on zero coupon yields. We follow Diebold et al. (2006) and compute

²¹ For Japan and Switzerland, the level slope and curvature factors only go up to October 21, 2013.

²² Note that this here is kurtosis and not excess kurtosis, and thus should be measured against a benchmark of 3. Nevertheless, this is still pretty high when compared to the other 4 currencies.

²³ The autocorrelation and partial-autocorrelation structure of the BDI returns (not shown here) in fact display the properties of an ARMA(1, 2) model.

²⁴ In a related context, high frequency yield curve data have been studied in Gürkaynak et al. (2005) and Brand et al. (2010) to assess the effect of central bank communication on various asset prices, including equities and exchange rates.

²⁵ More specifically, Diebold et al. (2006) show that their empirical level and slope factors have correlations of 43% and 39% with (year on year price deflator) inflation and capacity utilisation respectively (see page 319). Similarly, in a New Keynesian macro-finance model, Rudebusch and Wu (2008) find that their (macro-finance) level and slope factors have a 73% and 66% correlation with 1-year expected inflation and output (see page 916).

the three ‘empirical’ factors using linear combinations of yields of various maturities.²⁶ The level, slope and curvature factors are computed as:

$$\text{Level: } \mathcal{L}_t = (y_t^{(3)} + y_t^{(24)} + y_t^{(120)})/3 \quad (\text{A.1a})$$

$$\text{Slope: } \mathcal{S}_t = (y_t^{(3)} - y_t^{(120)}) \quad (\text{A.1b})$$

$$\text{Curvature: } \mathcal{C}_t = (2y_t^{(24)} - y_t^{(3)} - y_t^{(120)}), \quad (\text{A.1c})$$

where $y_t^{(\tau)}$ is the time t yield of a zero-coupon bond with maturity τ (measured in months). Zero coupon data for the US are taken from the well known and widely used [Gürkaynak et al. \(2007\)](#) database. For Australia, Canada and the UK, they are taken from the websites of the Reserve Bank of Australia (RBA), the Bank of Canada (BoC), and the Bank of England (BoE). For the euro area, the available yield curve data from the European Central Bank (ECB) only go back to the beginning of September 2004. To extend the data to the beginning of January 1999, we use yield curve factors from the Bundesbank before September 2004.²⁷ Due to the lack of publicly available daily data for Switzerland and Japan, we use the (daily) Nelson–Siegel–Svensson parameter estimates from [Malkhozov et al. \(2014\)](#) to construct the yield curve factors for these two countries.²⁸ The sample period that is covered by [Malkhozov et al. \(2014\)](#) nevertheless ends on the 21st of October 2013, thereby shortening the available out-of-sample evaluation period for the Yen and Swiss Franc somewhat.

To provide information about each country’s perceived macroeconomic fundamentals as captured by the yield curve factors relative to the US economy, we construct differences between the level, slope and curvature factors of the US and the foreign currency of interest. These are denoted by $\mathbf{x}_t^{\text{LSC},i}$ and computed as:

$$\mathbf{x}_t^{\text{LSC},i} = [(\mathcal{L}_t^{\text{US}} - \mathcal{L}_t^i), (\mathcal{S}_t^{\text{US}} - \mathcal{S}_t^i), (\mathcal{C}_t^{\text{US}} - \mathcal{C}_t^i)], \quad (\text{A.2})$$

where $\mathcal{L}_t^{\text{US}}$ (\mathcal{L}_t^i), $\mathcal{S}_t^{\text{US}}$ (\mathcal{S}_t^i), and $\mathcal{C}_t^{\text{US}}$ (\mathcal{C}_t^i) are level, slope and curvature factors for the US (the i th foreign currency), respectively, with $i = \{\text{EU, JP, GB, AU, CA, CH}\}$ being a country index for the exchange rates of interest. Due to the high persistence in the yield curve factors $\mathbf{x}_t^{\text{LSC},i}$, we use the (time) difference of the yield curve factors denoted by $\Delta \mathbf{x}_t^{\text{LSC},i}$, with Δ being the difference operator, as the predictor variables in the forecast evaluation.²⁹

A.2.2. Information in stock prices

We add stock price indices to the set of predictor variables to complement the information on macroeconomic fundamentals as contained in the yield curve. The usefulness of the information content embedded in US (as well as other countries’) stock returns for the purpose of forecasting the equity premium has recently been demonstrated by [Rapach et al. \(2013\)](#).³⁰ We use the SP500, as well as each

²⁶ Since it is common to use the Nelson–Siegel–Svensson approach to construct the zero coupon yields, as is done, for instance, by [Gürkaynak et al. \(2007\)](#), one could also use the slope coefficients fitted from the cross-sectional regression of the yields (ie., the β_i , $\forall i = 0, 1, 2$ estimates). Nevertheless, as is evident from the estimates that are provided in the Excel file provided by [Gürkaynak et al. \(2007\)](#) at <http://www.federalreserve.gov/pubs/feds/2006>, there can be considerable variation over time. We therefore prefer to construct the level, slope and curvature factors from the actual yield data.

²⁷ Note that both the ECB and the Bundesbank use the parametric approach of [Svensson \(1994\)](#) to construct zero coupon yields, so the methods of construction are consistent, despite the parameters being calibrated on two different sets of bonds. For the euro area, we use the [Svensson \(1994\)](#) parameter estimates reported under the ‘all issuers whose rating is triple A’ heading.

²⁸ We thank Andrea Vedolin for making these parameter estimates available to us. For more details on the construction of the factors and the bond data that were used, we refer to [Malkhozov et al. \(2014\)](#).

²⁹ Level, slope and curvature factors are known to be highly persistent, especially at daily frequencies. For the US, for instance, the first order autocorrelations are, 0.9994, 0.9986, and 0.9974 for $\mathcal{L}_t^{\text{US}}$, $\mathcal{S}_t^{\text{US}}$ and $\mathcal{C}_t^{\text{US}}$, respectively. Computing yield curve factors relative to the US ones remain highly persistent. Using euro area factors, this difference has autocorrelations of 0.9984, 0.9962, and 0.9914 for $(\mathcal{L}_t^{\text{US}} - \mathcal{L}_t^{\text{EU}})$, $(\mathcal{S}_t^{\text{US}} - \mathcal{S}_t^{\text{EU}})$ and $(\mathcal{C}_t^{\text{US}} - \mathcal{C}_t^{\text{EU}})$, respectively. To avoid this high persistence, we prefer to work with the (time) differenced series for $\mathbf{x}_t^{\text{LSC}}$, that is, $\Delta \mathbf{x}_t^{\text{LSC}}$.

³⁰ [Rapach et al. \(2013\)](#) show that lagged US returns have significant predictive power to forecast equity premia in numerous

individual country's headline stock price index, in the set of fundamental predictors. The headline indices are: the Nikkei225 for Japan, the FTSE100 for the UK, the SPI for Switzerland, the SPTSX for Canada and the All Ordinaries for Australia.³¹ For the euro area, it would seem natural to opt for the EURO STOXX 50 as a representative stock price index. Nevertheless, the key headline index from the viewpoint of the financial media still seems to be the DAX30, not only for Germany, but for the euro area as a whole. The DAX30 is also a more liquid market index. It has an approximately 50% higher trading volume (3 months average) than the EURO STOXX 50. For this reason, we prefer to use the DAX30 as the headline index for the euro area.

We construct returns of the stock price indices of the SP500 and each country's headline index to be used as predictor variables. That is, the stock (or equity) price predictor set (denoted by $\mathbf{x}_t^{\text{EQT},i}$) consists of:

$$\mathbf{x}_t^{\text{EQT},i} = [r_t^{\text{EQT,US}}, r_t^{\text{EQT},i}], \quad (\text{A.3})$$

where $r_t^{\text{EQT,US}}$ and $r_t^{\text{EQT},i}$ denote the return on the SP500 US stock price index and the return on the i th foreign equity market corresponding to the currency of interest, respectively. The i superscript here is again used to denote the foreign headline equity price index corresponding to the exchange rate of interest, i.e., $i = \{\text{DAX30, Nikkei225, FTSE100, AllOrds, SPTX, SPI}\}$.

A.2.3. Risk aversion measures and global trade activity

In addition to the yield curve and stock price data, we also include variables that are meant to capture risk aversion and global trade activity in the set of fundamental predictors.

We use the VIX index and the TED spread to provide us with a 'sense of risk aversion' in the market. The VIX measures the volatility implied by option prices on the SP500 and thus reflects investors' expectations about stock market volatility over the next month. The TED spread is calculated as the difference between the 3 month LIBOR rate (US dollar base) and the 3 month Treasury Bill rate and measures the perceived credit risk in the US economy. A higher value in the VIX and/or the TED spread is generally taken as an indication of market participants expecting an overall negative economic or financial outlook, and hence an increased (global) aversion to risk. Brunnermeier et al. (2009) have shown that the VIX and the TED spread predict higher returns in carry trade strategies which are widely used by foreign exchange traders.

We also include gold as a viable predictor variable. The motivation for this is twofold. First, gold is considered to be a 'safe haven' asset and hence constitutes a complement to the VIX and the TED spread indicators of risk aversion in financial markets. Gold is further regarded to be a hedge against inflation, deflation, as well as general uncertainties related to economic, financial and political instabilities. Second, together with other precious metals such as platinum and silver, gold is also commonly held in investment portfolios that are diversified over equities, bonds and exchange rates. Gold can thus be seen as a natural portfolio complement to foreign currency holdings in an investment portfolio. We expect, therefore, movements in gold prices to be informative for exchange rate forecasting, particularly since the financial crisis in 2008.

As a proxy for global trade flows as well as supply and demand trends in production of finished goods and raw materials, we include the Baltic Dry Index (BDI) and crude oil prices as fundamental predictor variables. The BDI is a composite index of the Baltic Capesize, Panamax, Handysize and Supramax indices. This index is designed as the successor to the Baltic Freight Index. The BDI is frequently viewed as a leading indicator of future global trade demand and economic growth, as the goods that are shipped are raw materials and thus give an indication of the demand for primary production

industrialised countries. The economic intuition behind this finding is given in relation to the US being an information originator (see pages 1635–1636 in Rapach et al., 2013 for a detailed explanation). Also, Hatemi-J et al. (2006) show that there can be causal predictive information in stock prices, which will be relevant also for exchange rate forecasts.

³¹ For Australia, we prefer to use the All Ordinaries index over the SP/ASX200 not only because of its longer data history, but also because it constitutes a broader index, containing 500 of the largest stocks as opposed to only 200 as the SP/ASX200.

inputs (see, for instance, Baumeister and Kilian, 2012, 2014; Baumeister et al., 2015 who also use the BDI as a measure of global trade activity to forecast oil prices).

The rationale for using crude oil prices in the set of predictors is due to oil still being one of the most widely used sources of energy (see for instance, among many other studies, the evidence reported in Lardic and Mignon, 2008 and He et al., 2010). Moreover, there is a widely held view that unexpected increases in the price of oil can cause recessions in many oil importing countries (see Kilian, 2008, Hamilton, 2009 and others). High oil prices are often also linked to periods of higher inflation, thereby directly affecting central bank policy and thus the setting of interest rates (Bhar and Mallik, 2013). Lastly, oil prices, in conjunction with US Energy Information Administration (EIA) inventories are closely monitored by financial market participants and reported in the financial press. These are taken to be early indicators of changes in production and manufacturing demand.

The risk and global trade activity predictor set, which we denote by $\mathbf{x}_t^{\text{RISK/ACTIV}}$, includes the following variables:

$$\mathbf{x}_t^{\text{RISK/ACTIV}} = [\Delta\text{VIX}_t, \Delta\text{TED}_t, r_t^{\text{GOLD}}, r_t^{\text{BDI}}, r_t^{\text{OIL}}], \quad (\text{A.4})$$

where ΔVIX_t , and ΔTED_t denote the (time) difference in the series of the CBOE Volatility Index and the TED spread, and $r_t^l, \forall l = \{\text{Gold, BDI, Oil}\}$ are the returns from investing in gold, the Baltic Dry Index and oil. Note here that we use the differences in the VIX and TED spread series. We could have also used the level series instead. Nevertheless, since both series are once again highly persistent, we have opted for the difference specification as used in Brunnermeier et al. (2009) as well.

All fundamental predictors which are used by fundamentalist agents to form their forecasts of currency i at time t are collected in the (10×1) dimensional vector:

$$[\mathbf{x}_t^{\text{LSC},i}, \mathbf{x}_t^{\text{EQT},i}, \mathbf{x}_t^{\text{RISK/ACTIV}}], \quad (\text{A.5})$$

where i denotes the foreign currency of interest.

A.2.4. Moving average rules

Moving average (MA) cross-over rules are among the most popular and common trading rules discussed in the technical analysis literature (Sullivan et al., 1999, p. 1656). The standard cross-over rule is that the down penetration of the MA by the price is regarded as a sell-signal and the upside penetration as a buy signal. There are various modifications to this rule. Buy and sell signals can be generated by crossovers of slow and fast MAs, where a slow MA is computed over a longer number of days than a fast MA.

Formally, the moving average of the exchange rate computed over the last n daily closing prices is defined as:

$$\text{MA}_t^{(n)} = n^{-1} \sum_{i=t-n+1}^t S_i, \quad \forall t \geq n \in \mathbb{Z}, \quad (\text{A.6})$$

where S_t is the daily spot closing price and n is the number of days that are averaged over. We consider the simplest and most widely used long-term cross-over of the 200 day moving average $\text{MA}_t^{(200)}$ and the spot price S_t .³²

³² Note that there are many other viable cross-over candidates involving cross-over rules of slow and fast moving MAs. These are generally the 50 and 100 day cross-overs with the 200 day MA. Nevertheless, to avoid any ambiguities related to 'searching over the best cross-over rule' issues, we stick to the simplest and most widely used long-term cross-over indicator, the crossing of the spot price S_t and the 200 day MA (see here also the relatively recent post on www.marketwatch.com with the title: What breaking the 200-day moving average for stocks really means for recent media coverage (article was published October 14, 2014) based on the spot price breaking through the 200 day MA).

The $MA_t^{(200)}$ is commonly viewed as a long-term trend indicator, with the indicator generating a broad buy signal as long as $S_t > MA_t^{(200)}$, while the penetration of the $MA_t^{(200)}$ by S_t from above reverses the signal to a sell indicator. Formally, we define the $S_t > MA_t^{(200)}$ buy indicator as:

$$IMA_t^{(200)} = \begin{cases} 1 & \text{if } S_t > MA_t^{(200)} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.7})$$

A.2.5. Momentum indicators

As an alternative to the moving average indicator, we also include a simple momentum indicator in the set of technical predictor variables. Momentum indicators are meant to capture the sentiment or trend following component in exchange rates, that is, the strategy to buy a currency if it had a positive return over the last n periods, and sell a currency if had a negative return. We use a time period of $n = 130$ days (6 months) to measure momentum, and define the 130 day momentum indicator as:

$$IMOM_t^{(130)} = \begin{cases} 1 & \text{if } S_t > S_{t-130} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.8})$$

The choice for 130 trading days (which corresponds to approximately half a year when 260 annual trading days are assumed) is mainly driven by a trade-off between the ability to capture known 'long-swings' in exchange rate data and to adapt quickly to recent changes. In the equity premium forecasting literature, it seems to be more common to use 9 months or 12 months horizons to compute the momentum indicator (see for instance page 4 in Neely et al., 2014). Nevertheless, the choice of using 6 months rather than 12 months returns to compute the momentum indicator does not have any important implications for our predictability results.³³

A.2.6. Relative strength index

We use the 14 day relative strength indices, denoted by $RSI_t^{(14)}$, in addition to the moving average and momentum indicators in the set of technical indicators. The RSI, as developed by Wilder (1978), measures the velocity of a security's price movement to identify overbought and oversold conditions. There exists recent empirical evidence illustrating the success of RSI based trading strategies. For instance, Chong and Ng (2008) use RSI based trading rules on the London Stock Exchange FT30 Index to analyse if these are profitable. Their conclusion is that an RSI based trading strategy is able to out-perform a simple buy-and-hold strategy. Similarly, Rodríguez-González et al. (2010) employ RSIs in a neural network context to predict individual stocks and are able to predict more than 50% of directions of change.

For a general n , the RSI is constructed as:

$$RSI_t^{(n)} = 100 - \frac{100}{1 + \frac{MA_t^{(n)}(dc_t)}{MA_t^{(n)}(uc_t)}} \quad (\text{A.9})$$

where $MA_t^{(n)}(\xi_t)$ denotes the n -period MA filter in (A.6) applied to variables ξ_t , and uc_t (dc_t) are upclose (downclose) measures defined as:

$$uc_t = \begin{cases} \Delta S_t & \text{if } \Delta S_t > 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad dc_t = \begin{cases} -\Delta S_t & \text{if } \Delta S_t < 0 \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A.10})$$

³³ Also, Neely et al. (2014) use volume data as a technical indicator. We do not do this here largely due to data availability. Volume data are much more difficult to get hold of, as exchange rates are still traded to a large extent over-the-counter.

with $\Delta S_t = S_t - S_{t-1}$ being the difference of the spot price of the exchange rate and n again the number of days over which the u_{c_t} and dc_t are averaged over. Note that RSIs are, by construction, an index over the 0 to 100 range.³⁴

To account for possible (traditional) time series dynamics in the returns, we also add lagged values of the returns to the set of technical predictors. The full set of technical predictor variables that is used by chartists to construct forecasts for the i th foreign currency is composed of the following 4 variables:³⁵

$$[r_t^i, \text{IMA}_t^{(200)}, \text{IMOM}_t^{(130)}, \text{RSI}_t^{(14)}]. \tag{A.11}$$

For reasons of completeness, we show the equivalent summary statistics corresponding to the technical indicators as used by chartist trading agents in Table A2. We group the three technical indicators of the 6 currencies into blocks for ease of readability. Note initially from the summary statistics that the moving average and momentum technical indicators, as defined in (A.7) and (A.8), are binary, so that the means measure the proportion of time that a buy signal was generated by the technicals. From these means one can notice that all indicators give readings of less than 0.5, suggesting that buy signals were generated less than 50% of the time over the approximately 15 years of data that are used in the forecast evaluation. The highest proportion of moving average based buy signals is generated for the British Pound, with the lowest being for the Australian Dollar. For the momentum based indicator, the largest proportion of buy signals is for the Yen, while the lowest is for the Swiss Franc. Overall, we can also notice from the first three ACFs and PACFs which are reported in the last six columns of Table A2 that the binary indicators for the moving average and momentum rules generate a fair degree of persistence in the predictor variables, with not only sizeable ACFs, but also PACFs.

The relative strength index based regressors, which are continuous but bounded series in the [0, 100] interval, are broadly centred at an RSI value of close to 50 for all 6 exchange rate series. Skewness and Kurtosis statistics indicate that the RSI based technical indicators are fairly symmetric without any showing of heavy tails. The ACFs and PACFs for the RSIs also indicate a noticeable degree of first order persistence in the series, similar to the persistence that an AR(1) process would generate.

A.3. Details of Dynamic Model Averaging (DMA) framework

To outline how the DMA framework is implemented, let y_t denote the variable to be predicted at time period t . Also, let \mathbf{x}_{t-1} be a $(1 \times K)$ vector that contains the full set of k predictors plus an intercept term ($K = k + 1$), and let $m = 1, \dots, M$ denote the model index, where $M = 2^k$ is the total number of possible (linear) model combinations (including the trivial model with only a constant term in it).³⁶ The set of predictors contained in the m th model is denoted by $\mathbf{x}_{t-1}^{(m)}$, with the dimension of $\mathbf{x}_{t-1}^{(m)}$ being $(1 \times K_m)$. The two equations that make up the DMA framework (for model m) are:

$$\text{Measurement: } y_t = \mathbf{x}_{t-1}^{(m)} \beta_t^{(m)} + u_t^{(m)} \tag{A.12a}$$

$(1 \times 1) \quad (1 \times K_m) \quad (K_m \times 1) \quad (1 \times 1)$

$$\text{State: } \beta_t^{(m)} = \beta_{t-1}^{(m)} + \epsilon_t^{(m)}, \tag{A.12b}$$

$(K_m \times 1) \quad (K_m \times 1) \quad (K_m \times 1)$

where (A.12a) and (A.12b) are measurement and state equations, respectively. The two disturbance terms $u_t^{(m)}$ and $\epsilon_t^{(m)}$ in (A.12) are jointly Multivariate Normal (MN) distributed, uncorrelated with each other and over time, that is:

³⁴ When working with stock price data, a stock is considered to be overbought when its RSI is above 70 and as oversold when its RSI is below 30. The choice of $n = 14$ is due to this being the most prominent value used among technical analysts, and in many software programmes is the default setting. Our results do not change in any important way if we use $n = 20$ instead, which is another popular setting.

³⁵ Note here, that, for simplicity of notation, we do not include i index counters on the technical indicators, but it should be clear that these are computed for the currency of interest.

³⁶ The term model here refers to the different possible linear combinations that can be obtained from using k predictors in a regression context, rather than the more general definition, where a model can be anything, potentially as flexible as non-linear or a non-parametric specification. The use of the term model is standard in the model averaging literature.

Table A2
Summary statistics of the technical predictor variables used by chartist forecasting agents.

Currency	Technical indicator	Mean	Median	Std. Dev.	Skew	Kurt	Min	Max	ACF(1)	ACF(2)	ACF(3)	PACF(1)	PACF(2)	PACF(3)
EUR	IMA _t ⁽²⁰⁰⁾	0.4108	0.0000	0.4920	0.3627	1.1315	0.0000	1.0000	0.9426	0.9214	0.9011	0.9427	0.2950	0.1066
	IMOM _t ⁽¹³⁰⁾	0.4150	0.0000	0.4928	0.3450	1.1191	0.0000	1.0000	0.9401	0.9220	0.9029	0.9404	0.3305	0.1172
	RSI _t ⁽¹⁴⁾	51.1616	50.7440	17.1477	0.0690	2.4367	1.1799	96.0865	0.9265	0.8569	0.7861	0.9269	-0.0119	-0.0477
JPY	IMA _t ⁽²⁰⁰⁾	0.4563	0.0000	0.4982	0.1754	1.0308	0.0000	1.0000	0.9190	0.8849	0.8602	0.9192	0.2602	0.1335
	IMOM _t ⁽¹³⁰⁾	0.4756	0.0000	0.4995	0.0976	1.0095	0.0000	1.0000	0.9214	0.8825	0.8515	0.9216	0.2227	0.1007
	RSI _t ⁽¹⁴⁾	49.8455	50.5988	16.9567	-0.0051	2.5552	0.1479	99.6377	0.9196	0.8442	0.7685	0.9199	-0.0091	-0.0437
GBP	IMA _t ⁽²⁰⁰⁾	0.4578	0.0000	0.4983	0.1694	1.0287	0.0000	1.0000	0.9404	0.9207	0.8961	0.9406	0.3158	0.0658
	IMOM _t ⁽¹³⁰⁾	0.4444	0.0000	0.4970	0.2236	1.0500	0.0000	1.0000	0.9317	0.9034	0.8811	0.9318	0.2692	0.1262
	RSI _t ⁽¹⁴⁾	51.1281	51.3032	16.9704	-0.0296	2.5395	1.4706	97.6821	0.9258	0.8533	0.7792	0.9266	-0.0287	-0.0495
AUD	IMA _t ⁽²⁰⁰⁾	0.3902	0.0000	0.4879	0.4500	1.2025	0.0000	1.0000	0.9302	0.8998	0.8778	0.9303	0.2577	0.1349
	IMOM _t ⁽¹³⁰⁾	0.4229	0.0000	0.4941	0.3121	1.0974	0.0000	1.0000	0.9197	0.8936	0.8640	0.9199	0.3079	0.0851
	RSI _t ⁽¹⁴⁾	52.3302	52.8121	16.5720	-0.0715	2.4626	0.7477	95.9903	0.9210	0.8449	0.7682	0.9210	-0.0217	-0.0453
CAD	IMA _t ⁽²⁰⁰⁾	0.4180	0.0000	0.4933	0.3326	1.1107	0.0000	1.0000	0.9138	0.8784	0.8532	0.9139	0.2642	0.1398
	IMOM _t ⁽¹³⁰⁾	0.4378	0.0000	0.4962	0.2509	1.0630	0.0000	1.0000	0.9142	0.8803	0.8609	0.9145	0.2705	0.1765
	RSI _t ⁽¹⁴⁾	51.7130	51.9048	16.1352	-0.0752	2.6570	0.4451	100.0000	0.9167	0.8383	0.7569	0.9179	-0.0116	-0.0603
CHF	IMA _t ⁽²⁰⁰⁾	0.3910	0.0000	0.4880	0.4468	1.1996	0.0000	1.0000	0.9271	0.8968	0.8686	0.9272	0.2662	0.0907
	IMOM _t ⁽¹³⁰⁾	0.4123	0.0000	0.4923	0.3564	1.1270	0.0000	1.0000	0.9202	0.8985	0.8749	0.9203	0.3390	0.1223
	RSI _t ⁽¹⁴⁾	51.0410	50.7385	17.1793	0.0726	2.5061	1.2376	99.2857	0.9251	0.8524	0.7786	0.9255	-0.0248	-0.0474

Notes: This table reports standard summary statistics of all the technical predictor variables used by chartist forecasting agents for exchange rate data from January 4, 1999 to June 30, 2014. The three technical indicators IMA_t⁽²⁰⁰⁾, IMOM_t⁽¹³⁰⁾ and RSI_t⁽¹⁴⁾ denote, respectively, the moving average indicator as defined in (A.7), the momentum indicator as defined in (A.8), and the RSI indicator as defined in (A.9). The abbreviation ACF(*p*) stands again for the *p*th autocorrelation function. We also include the partial autocorrelation function denoted by PACF(*q*) for order *q*th in this table to provide extra information about the dynamic structure of the technical indicators variables.

$$\begin{bmatrix} \mathbf{u}_t^{(m)} \\ \boldsymbol{\varepsilon}_t^{(m)} \end{bmatrix} \sim \text{MN} \left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} H_t^{(m)} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_t^{(m)} \end{bmatrix} \right), \tag{A.13}$$

where $H_t^{(m)}$ and $\mathbf{Q}_t^{(m)}$ are the variance and covariance matrix of the measurement and state equations, respectively.

Also, let \mathcal{M}_t denote the set of all possible models at time t , so that $\mathcal{M}_t \in \{1, 2, \dots, M\}$. Given knowledge of $H_t^{(m)}$ and $\mathbf{Q}_t^{(m)}$ and by fixing the model set $\mathcal{M}_t = m$, ie., to one particular model, the system in (A.12) takes the form of a standard state-space model, making it thereby possible to extract or ‘filter’ the time varying parameters $\boldsymbol{\beta}_t^{(m)}$ as the ‘latent states’ using standard Kalman Filter recursions. One-step ahead forecasts and forecast errors are available as a by-product of the Kalman Filter. Given $\mathcal{M}_t = m$, $H_t^{(m)}$ and $\mathbf{Q}_t^{(m)}$, the Kalman Filter recursions are:

$$\begin{aligned} \text{Prediction: } \hat{\boldsymbol{\beta}}_{t|t-1}^{(m)} &= \hat{\boldsymbol{\beta}}_{t-1|t-1}^{(m)} \\ \mathbf{P}_{t|t-1}^{(m)} &= \mathbf{P}_{t-1|t-1}^{(m)} + \mathbf{Q}_t^{(m)} \end{aligned} \tag{A.14a}$$

$$\hat{\mathbf{y}}_{t|t-1}^{(m)} = \mathbf{x}_{t-1}^{(m)} \hat{\boldsymbol{\beta}}_{t|t-1}^{(m)} \tag{A.14b}$$

$$\text{Prediction errors: } \hat{\mathbf{u}}_t^{(m)} = (\mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}^{(m)})$$

$$\text{MSE of prediction errors: } F_t^{(m)} = \mathbf{x}_{t-1}^{(m)} \mathbf{P}_{t|t-1}^{(m)} \mathbf{x}_{t-1}^{(m)\top} + H_t^{(m)} \tag{A.14c}$$

$$\text{Kalman Gain: } \mathbf{G}_t^{(m)} = \mathbf{P}_{t|t-1}^{(m)} \mathbf{x}_{t-1}^{(m)\top} / F_t^{(m)}$$

$$\text{Updating: } \hat{\boldsymbol{\beta}}_{t|t}^{(m)} = \hat{\boldsymbol{\beta}}_{t|t-1}^{(m)} + \mathbf{G}_t^{(m)} (\mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}^{(m)}) \tag{A.14d}$$

$$\mathbf{P}_{t|t}^{(m)} = \mathbf{P}_{t-1|t-1}^{(m)} - \mathbf{G}_t^{(m)} \mathbf{x}_{t-1}^{(m)} \mathbf{P}_{t-1|t-1}^{(m)}$$

where $\hat{\boldsymbol{\beta}}_{t|t-1}^{(m)} = \mathbb{E}_{t-1}(\boldsymbol{\beta}_t^{(m)})$, $\mathbb{E}_{t-1}(\cdot)$ is the expectation taken with respect to a time $t - 1$ information set denoted by \mathcal{I}_{t-1} , and $\mathbf{P}_{t|t-1}^{(m)}$ is the mean square error (MSE) of $\hat{\boldsymbol{\beta}}_{t|t-1}^{(m)}$. Forecasts from model m using information set \mathcal{I}_{t-1} are denoted by $\hat{\mathbf{y}}_{t|t-1}^{(m)}$. The one-step ahead forecast error is $\hat{\mathbf{u}}_t^{(m)}$ and its associated MSE is denoted by $F_t^{(m)}$. The $(K_m \times 1)$ vector $\mathbf{G}_t^{(m)}$ is the Kalman Gain. The terms $\hat{\boldsymbol{\beta}}_{t|t}^{(m)}$ and $\mathbf{P}_{t|t}^{(m)}$ are updated (or time t) estimates of the latent states $\boldsymbol{\beta}_t^{(m)}$ and their corresponding MSEs.

The Kalman Filter recursions in (A.14) are conditional on $H_t^{(m)}$ and $\mathbf{Q}_t^{(m)}$ (and model m). To avoid having to estimate $H_t^{(m)}$ and $\mathbf{Q}_t^{(m)}$, two simplifying assumptions are used in the literature. The first one, which is due to Raftery et al. (2010), is to replace $\mathbf{P}_{t|t-1}^{(m)}$ in (A.14a) by

$$\mathbf{P}_{t|t-1}^{(m)} = \frac{1}{\lambda} \mathbf{P}_{t-1|t-1}^{(m)}, \tag{A.15}$$

where $\lambda \in [0, 1]$. This approximation implies that $\mathbf{Q}_t^{(m)} = (\lambda^{-1} - 1) \mathbf{P}_{t-1|t-1}^{(m)}$. In the given context, the λ parameter is commonly referred to as a ‘forgetting factor’, as it determines how many observations are effectively used for estimation.³⁷ The second simplifying assumption is to replace the time varying

³⁷ This is also known as ‘windowing’. Intuitively, we can think of λ as a weighting function, where observations τ periods in the past receive a weight of λ^τ . See the discussion in Section 3.1 in Raftery et al. (2010) and pages 872–873 in Koop and Korobilis (2012) for more background and intuition about the use of forgetting factors in dynamic econometric models and what it implies for the effective sample size.

volatility $H_t^{(m)}$ by a simple exponentially weighted moving average (EWMA) estimate, that is, $H_t^{(m)}$ is constructed as:

$$H_t^{(m)} = \kappa H_{t-1}^{(m)} + (1 - \kappa) \hat{u}_{t-1}^{2(m)}, \tag{A.16}$$

where $\kappa \in [0, 1]$ is the standard EWMA smoothing parameter. Note here that an EWMA model can be thought of as a special form of a GARCH(1, 1) model, i.e., a restricted integrated GARCH(1, 1), with the restriction being that the intercept term is fixed at 0 and that the weights on the $t - 1$ volatility and squared error term sum to unity.³⁸

Model averaging or selection in the DMA framework is achieved by weighting the forecasts by their respective predictive model probabilities. To clarify this, let us define $\pi_{t|t-1}^{(m)}$ to be the probability of model m given information up to time $t - 1$, written as:

$$\pi_{t|t-1}^{(m)} = \Pr(\mathcal{M}_t = m | \mathcal{I}_{t-1}). \tag{A.17}$$

The DMA forecast of y_t , given information up to time $t - 1$, denoted as $\mathbb{E}(y_t | \mathcal{I}_{t-1})$, is then computed as:

$$\hat{y}_{t|t-1}^{(DMA)} = \sum_{m=1}^M \hat{y}_{t|t-1}^{(m)} \pi_{t|t-1}^{(m)}, \tag{A.18}$$

that is, as a weighted average of the forecasts from all possible models, $\{\hat{y}_{t|t-1}^{(m)}\}_{m=1}^M$, with the averaging weights being the predictive probabilities $\{\pi_{t|t-1}^{(m)}\}_{m=1}^M$.

To make the construction of the DMA forecasts in (A.18) feasible, model prediction and updating recursions are needed. Let $p_{jm} = \Pr(\mathcal{M}_t = m | \mathcal{M}_{t-1} = j)$ denote the (time invariant) transition probability of moving from model j at time $t - 1$ to model m at time t . Also, let $f_N^{(m)}(y_t | \mathcal{I}_{t-1})$ denote the predictive density of y_t given model m and information up to time $t - 1$. This predictive density is a Normal density evaluated at y_t with mean and variance given by $\hat{y}_{t|t-1}^{(m)}$ and $F_t^{(m)}$ as computed in (A.14b) and (A.14c), respectively. That is, $f_N^{(m)}(y_t | \mathcal{I}_{t-1}) = \mathcal{N}(\hat{y}_{t|t-1}^{(m)}, F_t^{(m)})$. Given an initial or prior model probability $\pi_{0|0}^{(m)}$, the model probability prediction and updating equations are then constructed as:

$$\text{Model Probability Prediction: } \pi_{t|t-1}^{(m)} = \sum_{j=1}^M \pi_{t-1|t-1}^{(j)} p_{jm} \tag{A.19a}$$

$$\text{Model Probability Updating: } \pi_{t|t}^{(m)} = \frac{\pi_{t-1|t-1}^{(m)} f_N^{(m)}(y_t | \mathcal{I}_{t-1})}{\sum_{j=1}^M \pi_{t-1|t-1}^{(j)} f_N^{(j)}(y_t | \mathcal{I}_{t-1})}. \tag{A.19b}$$

A final simplification that is needed to make the computation of the predictive model probabilities feasible is to approximate (A.19a) with

$$\pi_{t|t-1}^{(m)} = \frac{\pi_{t-1|t-1}^{\alpha(m)}}{\sum_{j=1}^M \pi_{t-1|t-1}^{\alpha(j)}}, \tag{A.20}$$

where $\alpha \in [0, 1]$. The approximation in (A.20) has the advantage that one avoids having to specify an $M \times M$ dimensional model probability transition matrix, which would make model prediction computationally infeasible when M is large. The α parameter in (A.20) can again be interpreted as a ‘forgetting factor’.

³⁸ It is well known in the volatility literature that GARCH(1, 1) models are difficult to beat in out-of-sample forecast evaluations (see, for instance, Hansen and Lunde, 2005). Approximating the time varying volatility by EWMA is thus unlikely to create any important loss in accuracy. We discuss later on how the κ parameter is calibrated.

The implementation of the DMA procedure to forecast exchange rate returns requires the calibration of the EWMA smoothing parameter κ , as well as the two forgetting factor parameters, λ and α . We follow the guidelines provided in RiskMetrics (1996, p. 97) for daily data and fix the κ parameter at 0.94. Koop and Korobilis (2012) recommend to set the values for λ and α close to 1, so that the parameters (as well as the model probabilities) evolve gradually over time.³⁹

A.4. Calibration of λ and α

We calibrate the λ and α parameters by minimising the mean squared forecast errors (MSE) via grid search on the first 500 observations, that is, on data from from January 6, 1999 to December 5, 2000. We use the following grid values for λ and α , respectively: $\alpha = [0.95 \ 0.99 \ 0.995 \ 0.999 \ 1]$ and $\lambda = [0.95 : 0.01 : 0.99 \ 0.995 : 0.001 : 0.999 \ 0.9995 \ 1]$. The notation $[a : c : b]$ is standard Matlab notation for: increment a by c up to b . Note here that performing a grid search over λ and α is computationally demanding. We therefore do this only once, ie., on the first 500 'in-sample' observations. The 'estimates' of λ and α are not updated as we roll through the sample. Also, due to the extra computational burden, we do not calibrate the κ parameter on the in-sample data, but simply fix κ at 0.94, which is the recommended value by RiskMetrics (1996) for daily data.

Notice from the λ and α grids that we specify a somewhat coarser grid for α , which is the forgetting factor used in the construction of the predictive model probabilities (see equation (A.20) in the model description section). As can be seen from the grid search plots below, there is not much variation in the MSE when it comes to different values of α . We limit our search over the $[0.95 \ 1]$ interval, which is the recommended range set by Koop and Korobilis (2012). For the λ parameter, there can be more variation at times and we thus find it important to use a somewhat finer search grid when considering values close to one.

In Figs. A2 and A3 we show the results from the grid search. In each figure, the vertical axis shows the MSE and the horizontal axis (the x -axis) the λ grid values. In each plot there are five differently coloured solid lines corresponding to the 5 α grid points that we use. The dotted horizontal and vertical lines mark the minimum MSE value and the corresponding λ value at minimum, respectively. The numeric λ and α values at min MSE after the grid search are also shown in the plots. From the plots in Figs. A2 and A3 it is clear that for all 6 exchange rate pairs, irrespective of whether fundamentalist or chartist predictor variables are used, the 'best' α value is attained at $\alpha^* = 0.99$. For 7 out of the 12 plots, the 'best' λ value is found at $\lambda^* = 0.999$. The remaining 5 are found at values close to 0.999, that is, two are at 0.9995 and 0.998 respectively, and one is 0.996. Given these grid search based optimal values, we calibrate for reasons of simplicity all fundamentalist and chartist DMA parameters at $\lambda = 0.999$ and $\alpha = 0.99$ and do not vary them for the 6 different exchange rates and the various forecast horizons that we consider.

A.5. Computing agent specific multiple period ahead out-of-sample forecasts

We also use the direct forecasting approach to construct individual, agent specific h -period ahead forecasts from the DMA framework. To do this, we re-write the relation in (A.12) (again using the general y_t and \mathbf{x}_t notation as

$$y_t = \mathbf{x}_{t-h}^{(m)} \boldsymbol{\beta}_{t,h}^{(m)} + u_t^{(m)} \quad (\text{A.21a})$$

$$\boldsymbol{\beta}_{t,h}^{(m)} = \boldsymbol{\beta}_{t-1,h}^{(m)} + \boldsymbol{\varepsilon}_t^{(m)}, \quad (\text{A.21b})$$

³⁹ See the discussion on pages 872–875 in Koop and Korobilis (2012). The effective window size, ie., how much of a weight observations in the past received, is determined from $1/(1-\lambda)$ (or $1/(1-\alpha)$, respectively). Choosing values below say 0.95, would make the window narrow, so that only the very recent past would receive non-zero weights, which could result in very noisy forecasts.

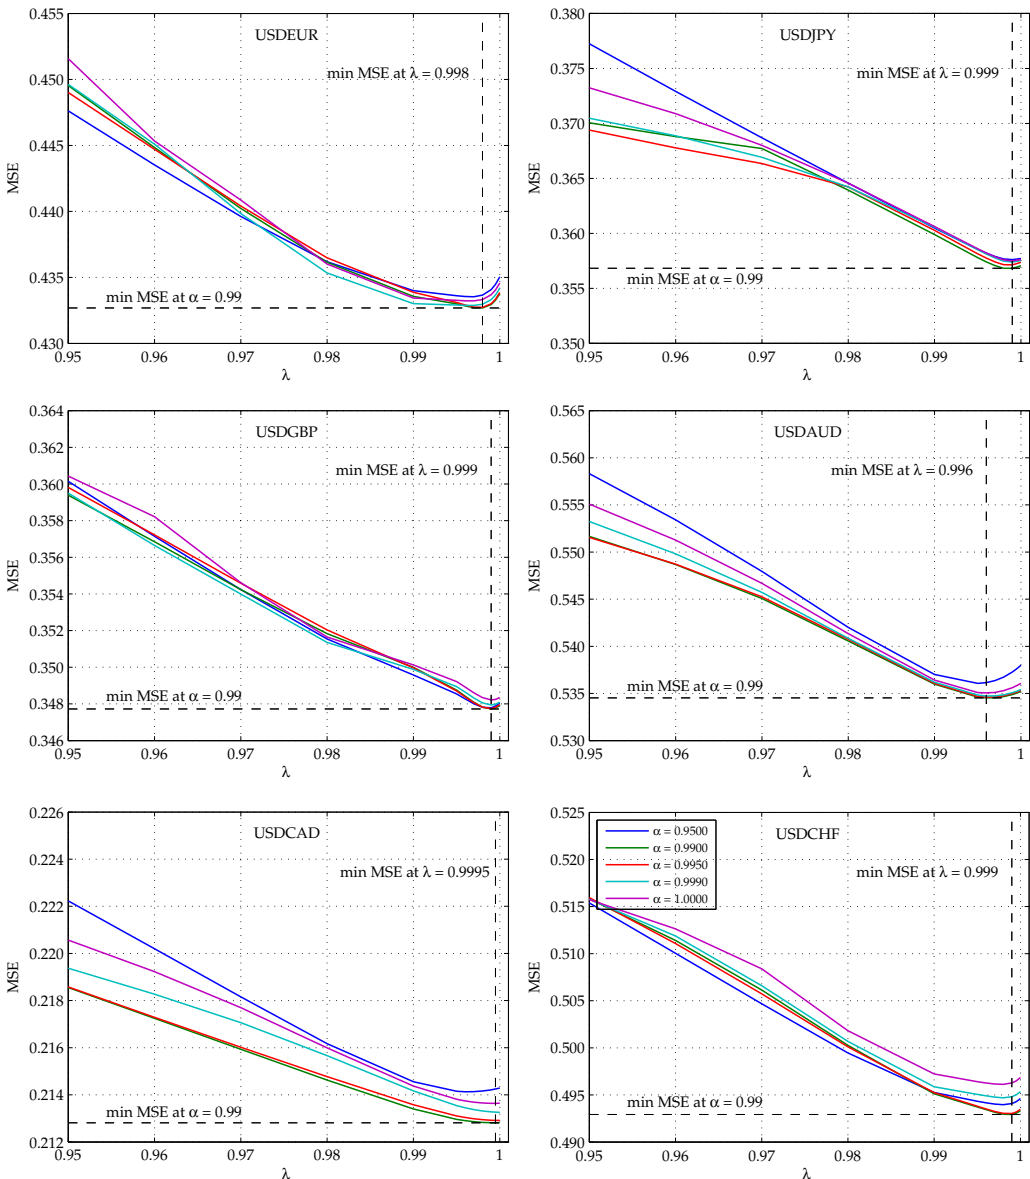


Fig. A2. Grid search over λ and α ($\kappa=0.94$): Fundamentalist model calibration results.

where the h subscript in $\beta_{t,h}^{(m)}$ signifies the relation to the h -period lagged value of \mathbf{x}_t . Using the same Kalman Filter recursions as in (A.14), but now on the h -period lagged relation as specified in (A.21) yields filtered estimates of the latent states (for each model and agent type $a = \{F, C\}$), that is, $\hat{\beta}_{t|t,h}^{A(m)}$. Given $\hat{\beta}_{t|t,h}^{A(m)}$, the DMA based agent specific h -step ahead forecasts are then computed as:

$$\hat{r}_{t+h|t}^A = \sum_{m=1}^M \mathbf{x}_t^{A(m)} \hat{\beta}_{t|t,h}^{A(m)} \pi_{t+h|t}^{A(m)} \tag{A.22a}$$

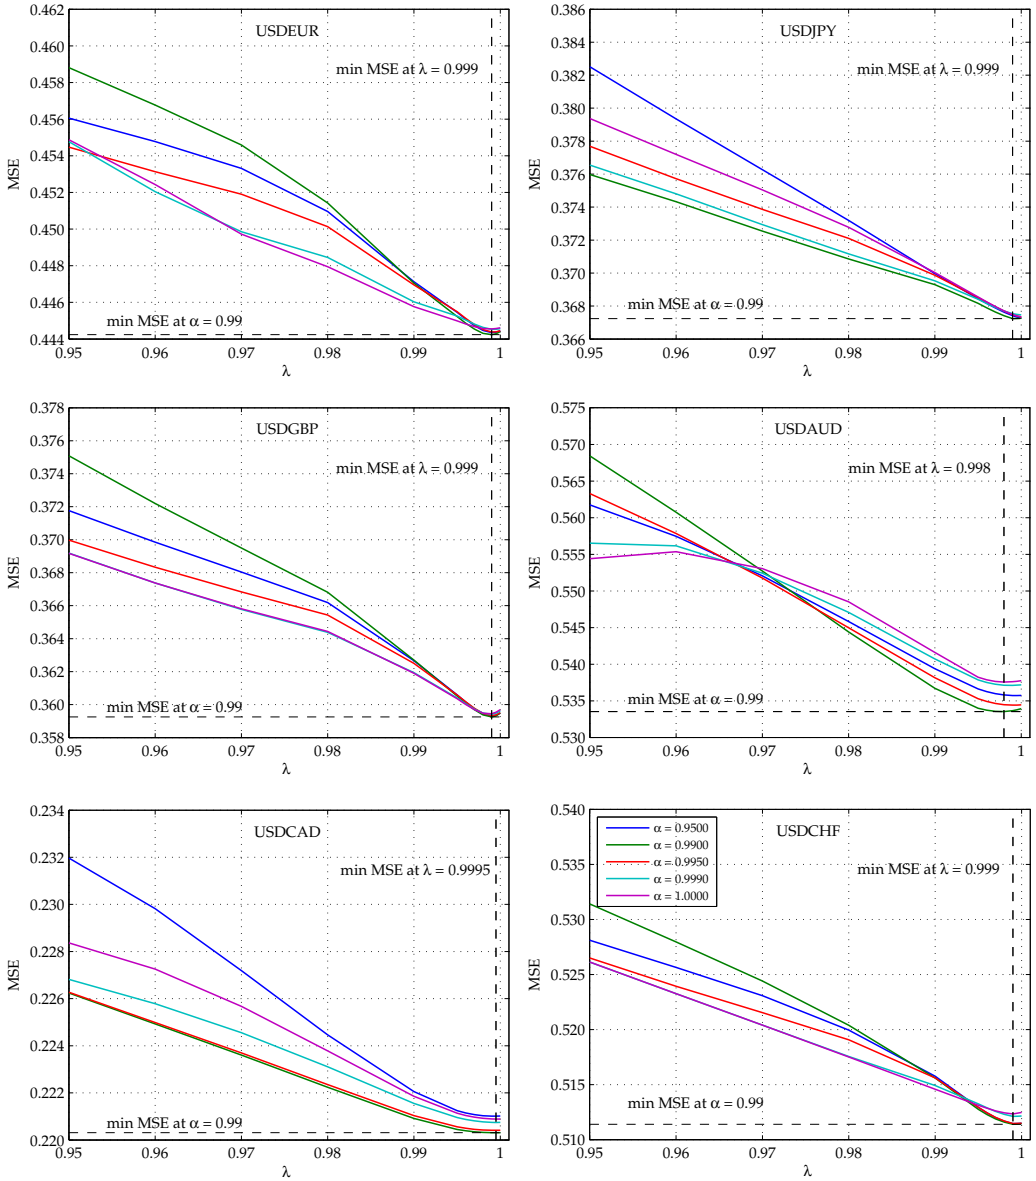


Fig. A3. Grid search over λ and α ($\kappa=0.94$): Chartist model calibration results.

for all $t = T_{is}, \dots, T$, where, again due to the random walk evolution of the latent state vector $\beta_{t+h,h}^{A(m)}$, the best forecast of $\beta_{t+h,h}^{A(m)}$ is its last observed filtered estimate, that is, $\mathbb{E}_t(\beta_{t+h,h}^{A(m)}) = \hat{\beta}_{t,t}^{A(m)}$. The h -step ahead agent specific predictive model probabilities at time t are computed from

$$\pi_{t+h|t}^{A(m)} = \frac{\pi_{t,t}^{A\alpha(m)}}{\sum_{j=1}^M \pi_{t,t}^{A\alpha(j)}}, \tag{A.23}$$

with $\pi_{it}^{A(m)}$ being the (agent specific) filtered model probability, analogous to the definition given in (A.19b), $\forall a = \{F, C\}$.

References

- Adrian, T., Etula, E., Shin, H.S., 2010. Risk appetite and exchange rates. Federal Reserve Bank of New York Staff Reports No. 361. Federal Reserve Bank of New York.
- Allen, H., Taylor, M.P., 1990. Charts, noise and fundamentals in the foreign exchange market. *Econ. J.* 100 (1), 4959.
- Andrews, D.W.K., Monahan, J.C., 1992. An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica* 60 (4), 953–966.
- Bates, J.M., Granger, C.W.J., 1969. The combination of forecasts. *J. Oper. Res. Soc.* 20 (4), 451–468.
- Baumeister, C., Kilian, L., 2012. Real-time forecasts of the real price of oil. *J. Bus. Econ. Stat.* 30 (2), 326–336.
- Baumeister, C., Kilian, L., 2014. What central bankers need to know about forecasting oil prices. *Int. Econ. Rev.* 55 (3), 869–889.
- Baumeister, C., Gurin, P., Kilian, L., 2015. Do high-frequency financial data help forecast oil prices? The MIDAS touch at work. *Int. J. Forecast.* 31 (2), 238–252.
- Bekaert, G., Hodrick, R.J., 1992. Characterizing predictable components in excess returns on equity and foreign exchange markets. *J. Finance* 47 (2), 467–509.
- Bhar, R., Mallik, G., 2013. Inflation uncertainty, growth uncertainty, oil prices, and output growth in the UK. *Empir. Econ.* 45 (3), 1333–1350.
- BIS, 2013. Foreign exchange turnover in April 2013: Preliminary global results. Triennial Central Bank Survey, Bank for International Settlements. <<http://www.bis.org/publ/rpfx13fx.pdf>> (accessed 06.14.).
- Brand, C., Buncic, D., Turunen, J., 2010. The impact of ECB monetary policy decisions and communication on the yield curve. *J. Eur. Econ. Assoc.* 8 (6), 1266–1298.
- Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. *Econometrica* 65 (5), 1059–1095.
- Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J. Econ. Dyn. Control* 22 (8–9), 1235–1274.
- Brunnermeier, M.K., Nagel, S., Pedersen, L.H., 2009. Carry trades and currency crashes. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), *NBER Macroeconomics Annual 2008*, vol. 23. University of Chicago Press, Chicago, IL, pp. 313–347s.
- Buncic, D., Gislser, K.I.M., 2015. Global equity market volatility spillovers: a broader role for the United States. Discussion Paper No. 2015–08, School of Economics and Political Sciences, University of St. Gallen. <<http://www1.vwa.unisg.ch/RePEc/usg/econwp/EWP-1508.pdf>> (accessed 06.14.).
- Buncic, D., Moretto, C., 2015. Forecasting copper prices with dynamic averaging and selection models. *N. Am. J. Econ. Finance* 33 (1), 1–38.
- Campbell, J.Y., Thompson, S.B., 2008. Predicting excess stock returns out of sample: can anything beat the historical average? *Rev. Financ. Stud.* 21 (4), 1509–1531.
- Cheung, Y.-W., Chinn, M.D., Pascual, A.G., 2003. What do we know about recent exchange rate models? In-sample fit and out-of-sample performance evaluated. Working Paper, Santa Cruz Center for International Economics, University of California, Santa Cruz.
- Chong, T.T.-L., Ng, W.-K., 2008. Technical analysis and the London stock exchange: testing the MACD and RSI rules using the FT30. *Appl. Econ. Lett.* 15 (14), 1111–1114.
- Claeskens, G., Magnus, J.R., Vasnev, A.L., Wang, W., 2015. The forecast combination puzzle: a simple theoretical explanation. SSRN Working Paper No. 2342841. <<http://dx.doi.org/10.2139/ssrn.2342841>> (accessed 06.14.).
- Clarida, R.H., Sarno, L., Taylor, M.P., Valente, G., 2003. The out-of-sample success of term structure models as exchange rate predictors: a step beyond. *J. Int. Econ.* 60 (1), 61–83.
- Clark, T.E., West, K.D., 2007. Approximately normal tests for equal predictive accuracy in nested models. *J. Econom.* 138 (1), 291–311.
- Croushore, D., Stark, T., 2001. A funny thing happened on the way to the data bank: a real time data set for macroeconomists. Fed. Reserve Bank of Philadelphia Bus. Rev. 5, 15–27.
- De Grauwe, P., Grimaldi, M., 2005. Heterogeneity of agents, transaction costs and the exchange rate. *J. Econ. Dyn. Control* 29 (4), 691–719.
- De Grauwe, P., Grimaldi, M., 2006. Exchange rate puzzles: a tale of switching attractors. *Eur. Econ. Rev.* 50 (1), 1–33.
- Della Corte, P., Tsiakas, I., 2012. Statistical and economic methods for evaluating exchange rate predictability. In: James, J., Marsh, I., Sarno, L. (Eds.), *Handbook of Exchange Rates*. Wiley, Hoboken, NJ, pp. 221–264.
- Della Corte, P., Sarno, L., Tsiakas, I., 2009. An economic evaluation of empirical exchange rate models. *Rev. Financ. Stud.* 22 (9), 3491–3530.
- Della Corte, P., Sarno, L., Tsiakas, I., 2011. Spot and forward volatility in foreign exchange. *J. Financ. Econ.* 100 (3), 496–513.
- Della Corte, P., Sarno, L., Sestieri, G., 2012. The predictive information content of external imbalances for exchange rate returns: how much is it worth? *Rev. Econ. Stat.* 94 (1), 100–115.
- Diebold, F.X., Rudebusch, G.D., Aruoba, S.B., 2006. The macroeconomy and the yield curve: a dynamic latent factor approach. *J. Econom.* 131 (1–2), 309–338.
- Elliott, G., Gargano, A., Timmermann, A., 2013. Complete subset regressions. *J. Econom.* 177 (2), 357–373.
- Goetzmann, W., Ingersoll, J., Spiegel, M., Welch, I., 2007. Portfolio performance manipulation and manipulation-proof performance measures. *Rev. Financ. Stud.* 20 (5), 1503–1546.
- Granger, C.W.J., Ramanathan, R., 1984. Improved methods of combining forecasts. *J. Forecast.* 3 (2), 197–204.
- Gürkaynak, R.S., Levin, A.T., Swanson, E.T., 2005. Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *Int. J. Cent. Bank.* 1 (1), 55–93.
- Gürkaynak, R.S., Sack, B., Wright, J.H., 2007. The U.S. treasury yield curve: 1961 to the present. *J. Monet. Econ.* 54 (8), 2291–2304.
- Hamilton, J.D., 2009. Causes and consequences of the oil shock of 2007–08. *Brookings Pap. Econ. Act.* 40 (1), 215–283.

- Hansen, P.R., Lunde, A., 2005. A forecast comparison of volatility models: does anything beat a GARCH(1,1)? *J. Appl. Econ. (Chichester Engl)* 20 (7), 873–889.
- Harvey, A.C., 1989. Forecasts of economic growth from the bond and stock markets. *Financ. Anal. J.* 45 (5), 38–45.
- Harvey, C.R., 1993. Term structure forecasts economic growth. *Financ. Anal. J.* 49 (3), 6–8.
- Hatemi-J, A., Roca, E., Buncic, D., 2006. Bootstrap causality tests of the relationship between the equity markets of the US and other developed countries: pre- and post-September 11. *J. Appl. Bus. Res.* 22 (3), 65–74.
- He, Y., Wang, S., Lai, K.K., 2010. Global economic activity and crude oil prices: a cointegration analysis. *Energy Econ.* 32 (4), 868–876.
- Huang, H., Lee, T.-H., 2010. To combine forecasts or to combine information? *Econom. Rev.* 29 (5–6), 534–570.
- Kilian, L., 2008. The economic effects of energy price shocks. *J. Econ. Lit.* 46 (4), 871–909.
- Koop, G., Korobilis, D., 2012. Forecasting inflation using dynamic model averaging. *Int. Econ. Rev.* 53 (3), 867–886.
- Koop, G., Korobilis, D., 2013. Large time-varying parameter VARs. *J. Econom.* 177 (2), 185–198.
- Lane, P.R., 2012. The European sovereign debt crisis. *J. Econ. Perspect.* 26 (3), 49–68.
- Lardic, S., Mignon, V., 2008. Oil prices and economic activity: an asymmetric cointegration approach. *Energy Econ.* 30 (3), 847–855.
- Li, J., Tsiakas, I., Wang, W., 2015. Predicting exchange rates out of sample: can economic fundamentals beat the random walk? *J. Financ. Econ.* 13 (2), 293–341.
- Lui, Y.-H., Mole, D., 1999. The use of fundamental and technical analysis by foreign exchange dealers: Hong Kong evidence. *J. Int. Money Finance* 17 (3), 535–545.
- Malkhozov, A., Mueller, P., Vedolin, A., Venter, G., 2014. Funding liquidity CAPM: international evidence. SSRN Working Paper No. 2420772. <<http://ssrn.com/abstract=2420772>> (accessed 06.14.).
- Mark, N.C., 1995. Exchange rates and fundamentals: evidence on long-horizon predictability. *Am. Econ. Rev.* 85 (1), 201–218.
- Meesse, R.A., Rogoff, K.S., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? *J. Int. Econ.* 14 (1–2), 3–24.
- Menkhoff, L., 1998. The noise trading approach questionnaire evidence from foreign exchange. *J. Int. Money Finance* 17 (3), 547–564.
- Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2013. Information flows in dark markets: dissecting customer currency trades. BIS Working Paper No. 405, Bank for International Settlements.
- Molodtsova, T., Papell, D.H., 2012. Taylor rule exchange rate forecasting during the financial crisis. NBER Working Paper No. 18330, National Bureau of Economic Research.
- Molodtsova, T., Nikolsko-Rzhevskyy, A., Papell, D.H., 2011. Taylor rules and the Euro. *J. Money Credit Bank* 43 (2–3), 535–552.
- Neely, C.J., Rapach, D.E., Tu, J., Zhou, G., 2014. Forecasting the equity risk premium: the role of technical indicators. *Manage. Sci.* 60 (7), 1772–1791.
- Raftery, A.E., Kárný, M., Ettler, P., 2010. Online prediction under model uncertainty via dynamic model averaging: application to a cold rolling mill. *Technometrics* 52 (1), 52–66.
- Rapach, D.E., Strauss, J.K., Zhou, G., 2010. Out-of-sample equity premium prediction: combination forecasts and links to the real economy. *Rev. Financ. Stud.* 23 (2), 821–862.
- Rapach, D.E., Strauss, J.K., Zhou, G., 2013. International stock return predictability: what is the role of the United States? *J. Finance* 68 (4), 1633–1662.
- Rime, D., Sarno, L., Sojli, E., 2010. Exchange rate forecasting, order flow and macroeconomic information. *J. Int. Econ.* 80 (1), 72–88.
- RiskMetrics, 1996. RiskMetrics: Technical Document, fourth ed. J. P. Morgan/Reuters. <<http://www.riskmetrics.com/system/files/private/td4e.pdf>> (accessed 06.14.).
- Rodríguez-González, A., Guldriés-Iglesias, F., Colomo-Palacios, R., et al., 2010. Improving trading systems using the RSI financial indicator and neural networks. In: Kang, B.-H., Richards, D. (Eds.), *Knowledge Management and Acquisition for Smart Systems and Services*. Springer-Verlag, Berlin Heidelberg, pp. 27–37.
- Rossi, B., 2013. Exchange rate predictability. *J. Econ. Lit.* 51 (4), 1063–1119.
- Rudebusch, G.D., Wu, T., 2008. A macro-finance model of the term structure, monetary policy and the economy. *Econ. J.* 118 (530), 906–926.
- Sarno, L., Taylor, M., 2003. *The Economics of Exchange Rates*. Cambridge University Press, Cambridge.
- Smith, J., Wallis, K.F., 2009. A simple explanation of the forecast combination puzzle. *Oxf. Bull. Econ. Stat.* 71 (3), 331–355.
- Sullivan, R., Timmermann, A., White, H., 1999. Data-snooping, technical trading rule performance, and the bootstrap. *J. Finance* 54 (5), 1647–1691.
- Svensson, L.E.O., 1994. Estimating and interpreting forward rates: Sweden 1992–4. NBER Working Paper No. 4871, National Bureau of Economic Research.
- Taylor, M.P., Allen, H., 1992. The use of technical analysis in the foreign exchange market. *J. Int. Money Finance* 11 (3), 304–314.
- Wilder, J.W., 1978. *New Concepts in Technical Trading Systems*. Trend Research, Greensboro, NC.
- Wright, J.H., 2008. Bayesian model averaging and exchange rate forecasts. *J. Econom.* 146 (2), 329–341.