DISCUSSION PAPER SERIES

No. 6598

AN ECONOMIC EVALUATION OF EMPIRICAL EXCHANGE RATE MODELS

Pasquale Della Corte, Lucio Sarno and Ilias Tsiakas

FINANCIAL ECONOMICS and INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

AN ECONOMIC EVALUATION OF EMPIRICAL EXCHANGE RATE MODELS

Pasquale Della Corte, University of Warwick
Lucio Sarno, University of Warwick, AXA Investment Managers and CEPR
Ilias Tsiakas, University of Warwick

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS and INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

ABSTRACT

An Economic Evaluation of Empirical Exchange Rate Models*

This paper provides a comprehensive evaluation of the short-horizon predictive ability of economic fundamentals and forward premia on monthly exchange rate returns in a framework that allows for volatility timing. We implement Bayesian methods for estimation and ranking of a set of empirical exchange rate models, and construct combined forecasts based on Bayesian Model Averaging. More importantly, we assess the economic value of the in-sample and out-of-sample forecasting power of the empirical models, and find two key results: (i) a risk averse investor will pay a high performance fee to switch from a dynamic portfolio strategy based on the random walk model to one which conditions on the forward premium with stochastic volatility innovations; and (ii) strategies based on combined forecasts yield large economic gains over the random walk benchmark. These two results are robust to reasonably high transaction costs.

JEL Classification: F31, F37 and G11
Keywords: Bayesian MCMC estimation, Bayesian model averaging, economic value, exchange rates, forward premium, monetary fundamentals and volatility

Pasquale Della Corte
Finance Group
Warwick Business School
University of Warwick
Coventry
CV4 7AL
United Kingdom
Email:
pasquale.dellacorte@wbs.ac.uk

Lucio Sarno
Finance Group
Warwick Business School
University of Warwick
Coventry
CV4 7AL
United Kingdom
Email: lucio.sarno@wbs.ac.uk

Ilias Tsiakas
Finance Group
Warwick Business School
University of Warwick
Coventry
CV4 7AL
United Kingdom
Email: ilias.tsiakas@wbs.ac.uk

Submitted 25 November 2007

Introduction

Forecasting exchange rates using models which condition on economically meaningful variables has long been at the top of the research agenda in international finance, and yet empirical success remains elusive. Starting with the seminal contribution of Meese and Rogoff (1983), a vast body of empirical research finds that models which condition on economic fundamentals cannot outperform a naive random walk model. Even though there is some evidence that exchange rates and fundamentals comove over long horizons (e.g. Mark, 1995; Mark and Sul, 2001), the prevailing view in international finance research is that exchange rates are not predictable, especially at short horizons.

A separate yet related literature finds that forward exchange rates contain valuable information for predicting spot exchange rates. In theory, the relation between spot and forward exchange rates is governed by the Uncovered Interest Parity (UIP) condition, which suggests that the forward premium must be perfectly positively related to future exchange rate changes. In practice, however, this is not the case as we empirically observe a negative relation.[1] The result of the empirical failure of UIP is that conditioning on the forward premium often generates exchange rate predictability. For example, Backus, Gregory and Telmer (1993) and Backus, Foresi and Telmer (2001) explore this further and find evidence of predictability using the lagged forward premium as a predictive variable. Furthermore, Clarida, Sarno, Taylor and Valente (2003, 2006) and Boudoukh, Richardson and Whitelaw (2006) show that the term structure of forward exchange (and interest) rates contains valuable information for forecasting spot exchange rates.

On the methodology side, while there is extensive literature on statistical measures of the accuracy of exchange rate forecasts, there is little work assessing the economic value of exchange rate predictability. Relevant research to date comprises an early study by West, Edison and Cho (1993) which provides a utility-based evaluation of exchange rate volatility, and more recently, Abhyankar, Sarno and Valente (2005) who use a similar method for investigating long-horizon exchange rate predictability. However, in the context of dynamic asset allocation strategies, there is no study assessing the economic value of the predictive ability of empirical exchange rate models which condition on economic fundamentals or the forward premium while allowing for volatility timing.

Our empirical investigation attempts to fill this gap and connect the related literatures which examine the performance of empirical exchange rate models. On the one hand, we directly investigate whether the statistical rejection of UIP generates economic value to a dynamically optimizing investor who exploits the UIP violation in order to generate excess returns. On the other hand, our economic evaluation provides a novel way for confirming (or not) the underwhelming performance of exchange rate models conditioning on economic fundamentals that has so far been established by statistical tests. We do this by employing a range of economic and Bayesian statistical criteria for performing a comprehensive assessment of the short-horizon, in-sample and out-of-sample, predictive ability of three sets of models for the conditional mean of monthly nominal exchange rate returns. These models include the naive random walk model, the monetary fundamentals model (in three variants), and the spot-forward regression model. Each of the models is studied under three volatility specifications: constant variance (standard linear regression), GARCH(1,1) and stochastic volatility (SV). In total, we evaluate the performance of 15 specifications, which encompass the most popular empirical exchange rate models studied in prior research. Our analysis employs monthly returns data ranging from January 1976 to December 2004 for three major US dollar exchange rates: the UK pound sterling, the Deutsche mark/euro, and the Japanese yen.

An important contribution of our analysis is the use of economic criteria. Statistical evidence

---

[1] See, for example, Bilson (1981), Fama (1984), Froot and Thaler (1990) and Backus, Foresi and Telmer (2001). For a survey of this literature, see Lewis (1995), Engel (1996) and the references therein.

of exchange rate predictability in itself does not guarantee that an investor can earn profits from an asset allocation strategy that exploits this predictability. In practice, ranking models is useful to an investor only if it leads to tangible economic gains. Therefore, we assess the economic value of exchange rate predictability by evaluating the impact of predictable changes in the conditional foreign exchange (FX) returns and volatility on the performance of dynamic allocation strategies. We employ mean-variance analysis as a standard measure of portfolio performance and apply quadratic utility, which allows us to quantify how risk aversion affects the economic value of predictability, building on empirical studies of volatility timing in stock returns by Fleming, Kirby and Ostdiek (2001) and Marquering and Verbeek (2004).[2] Ultimately, we measure how much a risk averse investor is willing to pay for switching from a dynamic portfolio strategy based on the random walk model to one which conditions on either monetary fundamentals or the forward premium and has a dynamic volatility specification.

The design of the dynamic allocation strategies is based on the mean-variance setting of West, Edison and Cho (1993), which adopts quadratic utility and fixes the investor's degree of relative risk aversion to a constant value. Quadratic utility allows us to compute in closed form the utility gains from using the conditional mean and volatility forecasts of one model rather than another. Combined with the approach of Fleming, Kirby and Ostdiek (2001), it is then straightforward to compute the performance fees, and hence provide an economically meaningful ranking of competing models for a given degree of relative risk aversion. Despite the well-known shortcomings affecting quadratic utility, there are a number of reasons that make it an appealing assumption, which we discuss in detail later in the paper. For example, quadratic utility is necessary to justify mean-variance optimization for non-normal return distributions, and therefore allows the economic evaluation of a larger universe of models within mean-variance. More importantly, there is evidence that quadratic utility provides a highly satisfactory approximation to a wide range of more sophisticated utility functions (e.g. Hlawitschka, 1994).

We assess the statistical evidence on exchange rate predictability in a Bayesian framework, which requires a choice for the prior distribution of the model parameters. For example, in the case of the simple linear regression model we assume independent Normal-Gamma prior distributions. We rank the competing model specifications by computing the posterior probability of each model. The posterior probability is based on the marginal likelihood and hence it accounts for parameter uncertainty, while imposing a penalty for lack of parsimony (higher dimension). In the context of this Bayesian methodology, an alternative approach to determining the best model available is to form combined forecasts which exploit information from the entire universe of model specifications under consideration. Specifically, we implement the Bayesian Model Averaging (BMA) method, which weighs all conditional mean and volatility forecasts by the posterior probability of each model.

To preview our key results, we find strong economic and statistical evidence against the naive random walk benchmark with constant variance innovations. In particular, while we confirm that conditioning on monetary fundamentals has no economic value either in-sample or out-of-sample, a key result of the paper is that the predictive ability of forward exchange rate premia has substantial economic value in a dynamic allocation strategy. Also, stochastic volatility significantly outperforms the constant variance and GARCH(1,1) models irrespective of the conditional mean specification. This leads to the conclusion that the best empirical exchange rate model is a model that exploits the information in the forward market for the prediction of conditional exchange rate returns and allows for stochastic volatility for the prediction of exchange rate volatility. We also provide evidence

---

[2]For studies of asset return predictability following this approach see also Kandell and Stambaugh (1996), Barberis (2000), Baks, Metrick and Wachter (2001), Bauer (2001), Shanken and Tamayo (2001), Avramov (2002), Cremers (2002) and Della Corte, Sarno and Thornton (2007). Karolyi and Stulz (2003) provide a survey of asset allocation in an international context.

that combined forecasts which are formed using BMA substantially outperform the random walk benchmark. These results are robust to reasonably high transaction costs and hold for all currencies both in-sample and out-of-sample. Finally, these findings have clear implications for international asset allocation strategies which are subject to FX risk.

The remainder of the paper is organized as follows. In the next section we briefly review the literature on exchange rate predictability conditioning on either fundamentals or forward exchange premia. Section 2 lays out the competing empirical models for the conditional mean and volatility of exchange rate returns. In Section 3 we discuss the framework for assessing the economic value of exchange rate predictability for a risk averse investor with a dynamic portfolio allocation strategy. Section 4 provides a sketch of the Bayesian estimation tools, discusses the approach to model selection, and explains the construction of combined forecasts using the BMA method. Our empirical results are reported in Section 5, followed by robustness checks in Section 6. Finally, Section 7 concludes.

1 Stylized Facts on Exchange Rate Predictability

In this section we briefly review the theoretical and empirical research that motivates our conditioning on lagged monetary fundamentals and forward premia in the set of empirical exchange rate models.

1.1 Exchange Rates and Monetary Fundamentals

There is extensive literature in international finance which studies the relation between nominal exchange rates and monetary fundamentals and focuses on the following predictive variable, x_t:

$$x_t = z_t - s_t \tag{1}$$

$$z_t = (m_t - m_t^*) - (y_t - y_t^*) \tag{2}$$

where s_t is the log of the nominal exchange rate (defined as the domestic price of foreign currency); m_t is the log of the money supply; y_t is the log of national income; and asterisks denote variables of the foreign country. Note that long-run money neutrality and income homogeneity are imposed, with the coefficients on $m_t - m_t^*$ and $y_t - y_t^*$ both set to unity as predicted by conventional theories of exchange rate determination, and z_t represents the relative velocity between the two countries in question. The relation between exchange rates and fundamentals defined in Equations 1 and 2 suggests that a deviation of the nominal exchange rate, s_t, from its long-run equilibrium level determined by the fundamentals, z_t (i.e. $x_t \neq 0$), requires the exchange rate to move in the future so as to converge towards its long-run equilibrium. In other words, the deviation x_t has predictive power on future realizations of the exchange rate.[3]

Despite the appeal of the theoretical relation between exchange rates and fundamentals, the empirical evidence is mixed. On the one hand, short-run exchange rate variability appears to be disconnected from the underlying fundamentals (Mark, 1995) in what is commonly referred to as the "exchange rate disconnect puzzle". On the other hand, some recent empirical research finds that fundamentals and nominal exchange rates move together *in the long run* (Groen, 2000; Mark and Sul, 2001; Rapach and Wohar, 2002; Sarno, Valente and Wohar, 2004). Either way, our study contributes to the empirical literature on the predictive ability of monetary fundamentals on exchange

---

[3]The specification of fundamentals in Equation 2 is common in the relevant empirical literature; e.g. see Mark (1995) and Mark and Sul (2001). Theories of exchange rate determination view z_t as the core set of economic fundamentals that determine the long-run equilibrium exchange rate. These theories include traditional models based on aggregate demand functions (e.g. Mark, 1995; Engel and West, 2005), and representative-agent general equilibrium models (e.g. Lucas, 1982; Obstfeld and Rogoff, 1995).

rates by providing an economic evaluation of the in-sample and out-of-sample forecasting power of fundamentals at a *short* (one-month ahead) horizon.

1.2 The Spot-Forward Exchange Rate Relation

Assuming risk neutrality and rational expectations, Uncovered Interest Parity (UIP) is the cornerstone condition for FX market efficiency. For a one-period horizon, UIP is represented by the following equation:

$$E_{t-1}\Delta s_t = i_{t-1} - i_{t-1}^* \tag{3}$$

where i_{t-1} and i_{t-1}^* are the one-period domestic and foreign nominal interest rates respectively; and $\Delta s_t \equiv s_t - s_{t-1}$.

In the absence of riskless arbitrage, Covered Interest Parity (CIP) holds and implies:

$$f_{t-1} - s_{t-1} = i_{t-1} - i_{t-1}^* \tag{4}$$

where f_{t-1} is the log of the one-period forward exchange rate (i.e. the rate agreed now for an exchange of currencies in one period). Substituting the interest rate differential $i_{t-1} - i_{t-1}^*$ in Equation 3 by the forward premium (or forward discount) $f_{t-1} - s_{t-1}$, we can estimate the following regression, which is commonly referred to as the "Fama regression" (Fama, 1984):

$$\Delta s_t = \alpha + \beta \left(f_{t-1} - s_{t-1} \right) + u_t \tag{5}$$

where u_t is a disturbance term.

If UIP holds, we should find that $\alpha = 0$, $\beta = 1$, and the disturbance term u_t is uncorrelated with information available at time $t-1$. Despite the increasing sophistication of the econometric techniques implemented and the improving quality of the data sets utilized, empirical studies estimating the Fama regression consistently reject the UIP condition (Hodrick, 1987; Lewis, 1995; Engel, 1996). As a result, it is now a stylized fact that estimates of β tend to be closer to minus unity than plus unity (Froot and Thaler, 1990). The negative value of β is the defining feature of what is commonly referred to as the "forward bias puzzle," namely the tendency of high-interest currencies to appreciate when UIP would predict them to depreciate.[4]

Attempts to explain the forward bias puzzle using models of risk premia have met with limited or mixed success, especially for plausible degrees of risk aversion (e.g. Engel, 1996, and the references therein). Moreover, it has proved difficult to explain the rejection of UIP by resorting to a range of proposed explanations, including learning, peso problems and bubbles (e.g. Lewis, 1995); consumption-based asset pricing theories, which allow for departures from both time-additive preferences (Backus, Gregory and Telmer, 1993; Bansal, Gallant, Hussey and Tauchen, 1995; Bekaert, 1996) and from expected utility (Bekaert, Hodrick and Marshall, 1997); and using popular models of the term structure of interest rates adapted to a multi-currency setting (Backus, Foresi and Telmer, 2001). In conclusion, even with the benefit of twenty years of hindsight, the forward bias has not been convincingly explained and remains a puzzle in international finance research.

In this context, the objective of this paper is neither to find a novel resolution to the forward bias puzzle nor to discriminate among competing explanations. Instead, we focus on predicting short-horizon exchange rate returns when conditioning on the lagged forward premium, thus empirically

---

[4]Exceptions to this puzzle include Bansal (1997), who finds that the forward bias is related to the sign of the interest rate differential; Bansal and Dahlquist (2000), who document that the forward bias is largely confined to developed economies and countries where the interest rate is lower than the US; Bekaert and Hodrick (2001), who provide a "partial rehabilitation" of UIP by accounting for small-sample distortions; and Lustig and Verdelhan (2007), who attempt to explain the forward bias puzzle focusing on the cross-section of foreign currency risk premia.

exploiting the forward bias reported in the strand of literature stemming from Bilson (1981), Fama (1984), Bekaert and Hodrick (1993) and Backus, Gregory and Telmer (1993). For example, Bilson (1981) argues that regressions conditioning on the forward premium can potentially yield substantial economic returns, whereas arguments based on limits to speculation would suggest otherwise (Lyons, 2001; Sarno, Valente and Leon, 2006). Furthermore, term structure models that exploit departures from UIP often yield accurate out-of-sample forecasts (e.g. Clarida and Taylor, 1997; Clarida, Sarno, Taylor and Valente, 2003; Boudoukh, Richardson and Whitelaw, 2006). However, little attention has been given to the question of whether the statistical rejection of UIP and the forward bias resulting from the negative estimate of β offers economic value to an international investor facing FX risk. Our paper fills this void in the literature by assessing the economic value of the predictive ability of empirical exchange rate models which condition on the forward premium in the context of dynamic asset allocation strategies.

2 Modeling FX Returns and Volatility

In this section we present the candidate models applied to monthly exchange rate returns in our study of short-horizon exchange rate predictability. We use a set of specifications for the dynamics of both the conditional mean and volatility, which are set against the naive random walk benchmark. In short, we estimate five conditional mean and three conditional volatility specifications yielding a total of 15 models for each of the three dollar exchange rates under consideration.

2.1 The Conditional Mean

We examine five conditional mean specifications in which the dynamics of exchange rate returns are driven by the following regression:

$$\Delta s_t = \alpha + \beta x_{t-1} + u_t, \quad u_t = v_t \varepsilon_t, \quad \varepsilon_t \sim NID\,(0,1)\,. \tag{6}$$

Our first specification is the naive random walk (RW) model, which sets $\beta = 0$. This model is the standard benchmark in the literature on exchange rate predictability since the seminal work of Meese and Rogoff (1983).

The next three model specifications condition on monetary fundamentals (MF). Specifically, MF_1 uses the canonical version $x_t = z_t - s_t$ as defined in Equations 1 and 2. This is the most common formulation of the monetary fundamentals model since Mark (1995). The second variant, MF_2, corrects for the deterministic component in the deviation of the exchange rate from monetary fundamentals by allowing for an intercept and a slope parameter; in other words, we run the ordinary least squares (OLS) regression $s_t = \kappa_0 + \kappa_1 z_t + \zeta_t$, and set $x_t = -\widehat{\zeta}_t$. The third variant, MF_3, further corrects for the time trend in fundamentals deviations; in this case, we run the OLS regression $s_t = \kappa_0 + \kappa_1 z_t + \kappa_2 t + \zeta_t$, where t is a simple time trend, and again we set $x_t = -\widehat{\zeta}_t$.[5]

Finally, the fifth conditional mean specification is the forward premium (FP) model, which sets $x_t = f_t - s_t$ as in Equation 5 resulting in the Fama (1984) regression. The FP model stems directly from the spot-forward exchange rate relation derived from UIP. Hence it constitutes the empirical model which exploits the forward bias and allows us to assess the economic value of conditioning on the forward premium in the context of dynamic asset allocation strategies. The forward bias (a

---

[5]The motivation behind the MF_2 and MF_3 variants derives from empirical evidence that cointegration between s_t and z_t may be established only by correcting for the deterministic components (either a constant or a constant and a time trend) in the cointegrating residual (e.g. Rapach and Wohar, 2002). Note, however, that in the out-of-sample exercise we estimate the deterministic component recursively as we move through the data sample, and hence our results do not suffer from "look-ahead bias".

negative estimate of the β coefficient in the FP model) implies that the more the foreign currency is at a premium in the forward market, the less the home currency is expected to depreciate. Equivalently, the more domestic interest rates exceed foreign interest rates, the more the domestic currency tends to appreciate over the holding period.

2.2 The Conditional Variance

We model the dynamics of the conditional variance by implementing three models: the simple linear regression (LR), the GARCH(1,1) model, and the stochastic volatility (SV) model. The linear regression framework simply assumes that the conditional variance of FX return innovations is constant over time $\left(v_t^2 = v^2\right)$, and therefore presents the benchmark against which models with time-varying conditional variance will be evaluated.

The benchmark GARCH(1,1) model of Bollerslev (1986) is defined as:

$$v_t^2 = \omega + \gamma_1 v_{t-1}^2 + \gamma_2 u_{t-1}^2. \tag{7}$$

Our motivation for studying the simple GARCH(1,1) model is based on the early study of West, Edison and Cho (1993), which conducts a utility-based evaluation of exchange rate volatility and finds that GARCH(1,1) is the best performing model.

Stochastic volatility models are similar to the GARCH process in that they capture the persistent and hence predictable component of volatility. Unlike GARCH models, however, the assumption of a stochastic second moment introduces an additional source of risk.[6] According to the plain vanilla SV model, the persistence of the conditional volatility v_t is captured by the dynamics of the Gaussian stochastic log-variance process h_t:

$$v_t = \exp\left(h_t/2\right) \tag{8}$$

$$h_t = \mu + \phi\left(h_{t-1} - \mu\right) + \sigma\eta_t, \quad \eta_t \sim NID\left(0,1\right). \tag{9}$$

3 Measuring the Economic Value of Exchange Rate Predictability

This section discusses the framework we use in order to evaluate the impact of predictable changes in both exchange rate returns and volatility on the performance of dynamic allocation strategies.

3.1 FX Models in a Dynamic Mean-Variance Framework

In mean-variance analysis, the maximum expected return strategy leads to a portfolio allocation on the efficient frontier. Consider an investor who has a one-month horizon and constructs a dynamically rebalanced portfolio that maximizes the conditional expected return subject to achieving a target conditional volatility. Computing the time-varying weights of this portfolio requires one-step ahead forecasts of the conditional mean and the conditional variance-covariance matrix. Let r_{t+1} denote the $K \times 1$ vector of risky asset returns; $\mu_{t+1|t} = E_t\left[r_{t+1}\right]$ is the conditional expectation of r_{t+1}; and $\Sigma_{t+1|t} = E_t\left[\left(r_{t+1} - \mu_{t+1|t}\right)\left(r_{t+1} - \mu_{t+1|t}\right)'\right]$ is the conditional variance-covariance matrix of r_{t+1}.

---

[6]Market microstructure theories of speculative trading (e.g. Tauchen and Pitts, 1983; Andersen, 1996) provide rigorous arguments for modeling volatility as stochastic. For details on SV models see Kim, Shephard and Chib (1998) and Chib, Nardari and Shephard (2002). For an application of SV models to exchange rates, see Harvey, Ruiz and Shephard (1994). Finally, for a comparison of GARCH and SV models see Fleming and Kirby (2003).

At each period t, the investor solves the following problem:

$$\max_{w_t} \left\{ \mu_{p,t+1|t} = w_t'\mu_{t+1|t} + \left(1 - w_t'\iota\right) r_f \right\}$$

$$\text{s.t.} \quad \left(\sigma_p^*\right)^2 = w_t'\Sigma_{t+1|t}w_t \tag{10}$$

where w_t is the $K \times 1$ vector of portfolio weights on the risky assets; ι is a $K \times 1$ vector of ones; $\mu_{p,t+1|t}$ is the conditional expected return of the portfolio; σ_p^* is the target conditional volatility of the portfolio returns; and r_f is the return on the riskless asset. The solution to this optimization problem delivers the risky asset weights:

$$w_t = \frac{\sigma_p^*}{\sqrt{C_t}}\Sigma_{t+1|t}^{-1}\left(\mu_{t+1|t} - \iota r_f\right) \tag{11}$$

where $C_t = \left(\mu_{t+1|t} - \iota r_f\right)'\Sigma_{t+1|t}^{-1}\left(\mu_{t+1|t} - \iota r_f\right)$. The weight on the riskless asset is $1 - w_t'\iota$.

Constructing the optimal portfolio weights requires estimates of the conditional expected returns, variances and covariances. We consider five conditional mean strategies (RW, MF_1, MF_2, MF_3 and FP) and three conditional volatility strategies (LR, $GARCH$ and SV) for a total of 15 sets of one-step ahead conditional expected return and volatility forecasts. The conditional covariances are computed using the constant conditional correlation (CCC) model of Bollerslev (1990), in which the dynamics of covariances are driven by the time-variation in the conditional volatilities.[7] By design, in this setting the optimal weights will vary across models only to the extent that forecasts of the conditional mean and volatility will vary, which is precisely what the empirical models provide. The benchmark against which we compare the model specifications is the random walk model with constant variance (RW^{LR}). In short, our objective is to determine whether there is economic value in (i) conditioning on monetary fundamentals and, if so, which of the three specifications works best, (ii) conditioning on the forward premium, (iii) using a GARCH volatility specification, and (iv) implementing an SV process for the monthly FX innovations.

3.2 Quadratic Utility

Mean-variance analysis is a natural framework for assessing the economic value of strategies which exploit predictability in the mean and variance. In particular, we rank the performance of the competing FX models using the West, Edison and Cho (1993) methodology, which is based on mean-variance analysis with quadratic utility. The investor's realized utility in period $t+1$ can be written as:

$$U\left(W_{t+1}\right) = W_{t+1} - \frac{\lambda}{2}W_{t+1}^2 = W_t R_{p,t+1} - \frac{\lambda W_t^2}{2}R_{p,t+1}^2 \tag{12}$$

where W_{t+1} is the investor's wealth at $t+1$, λ determines his risk preference, and

$$R_{p,t+1} = 1 + r_{p,t+1} = 1 + \left(1 - w_t'\iota\right) r_f + w_t'r_{t+1} \tag{13}$$

is the period $t+1$ gross return on his portfolio.

We quantify the economic value of exchange rate predictability by setting the investor's degree of relative risk aversion (RRA) $\delta_t = \lambda W_t/\left(1 - \lambda W_t\right)$ equal to a constant value δ. In this case, West,

---

[7]In notation local to this footnote, the CCC model of Bollerslev (1990) specifies the covariances as follows: $\sigma_{ij,t} = \sigma_{i,t}\sigma_{j,t}\rho_{ij}$, where $\sigma_{i,t}$ and $\sigma_{j,t}$ are the conditional volatilities implied by either the GARCH(1,1) or the SV process, and ρ_{ij} is the constant sample correlation coefficient. Note that for the out-of-sample results we use a rolling correlation estimate updated every time a new observation is added. From a numerical standpoint, implementing the CCC model is attractive because it eliminates the possibility of $\Sigma_{t+1|t}$ not being positive definite.

Edison and Cho (1993) demonstrate that one can use the average realized utility, $\overline{U}(\cdot)$, to consistently estimate the expected utility generated by a given level of initial wealth. Specifically, the average utility for an investor with initial wealth W_0 is equal to:

$$\overline{U}(\cdot) = W_0 \sum_{t=0}^{T-1} \left\{ R_{p,t+1} - \frac{\delta}{2(1+\delta)} R_{p,t+1}^2 \right\}. \tag{14}$$

Average utility depends on taste for risk. In the absence of restrictions on δ, quadratic utility exhibits increasing RRA. This is counterintuitive since, for instance, an investor with increasing RRA becomes more averse to a percentage loss in wealth when his wealth increases. As in West, Edison and Cho (1993) and Fleming, Kirby and Ostdiek (2001), fixing δ implies that expected utility is linearly homogeneous in wealth: double wealth and expected utility doubles. Hence we can standardize the investor problem by assuming $W_0 = \$1$. Furthermore, by fixing δ rather than λ, we are implicitly interpreting quadratic utility as an approximation to a non-quadratic utility function, with the approximating choice of λ dependent on wealth. The estimate of expected quadratic utility given in Equation 14 is used to implement the Fleming, Kirby and Ostdiek (2001) framework for assessing the economic value of our FX strategies in the context of dynamic asset allocation.

A critical aspect of mean-variance analysis is that it applies exactly only when the return distribution is normal *or* the utility function is quadratic. Hence, the use of quadratic utility is not necessary to justify mean-variance optimization. For instance, one could instead consider using utility functions belonging to the constant relative risk aversion (CRRA) class, such as power or log utility. However, quadratic utility is an attractive assumption because it allows us to consider non-normal distributions of returns, while remaining within the mean-variance framework as well as providing a high degree of analytical tractability. Absent Gaussianity, quadratic utility is needed to justify mean-variance and allows us to use the Fleming, Kirby and Ostdiek (2001) framework (also based on quadratic utility) for evaluating the performance of fat-tailed volatility specifications, such as the $tGARCH$ model of Bollerslev (1987).

Additionally, quadratic utility may be viewed as a second order Taylor series approximation to expected utility. In an investigation of the empirical robustness of the quadratic approximation, Hlawitschka (1994) finds that a two-moment Taylor series expansion "may provide an excellent approximation" (p. 713) to expected utility and concludes that the ranking of common stock portfolios based on two-moment Taylor series is "almost exactly the same" (p. 714) as the ranking based on a wide range of utility functions.

3.3   Performance Measures

At any point in time, one set of estimates of the conditional mean and variance is better than a second set if investment decisions based on the first set lead to higher average realized utility, $\overline{U}$. Alternatively, the optimal model requires less wealth to yield a given level of $\overline{U}$ than a suboptimal model. Following Fleming, Kirby and Ostdiek (2001) we measure the economic value of our FX strategies by equating the average utilities for selected pairs of portfolios. Suppose, for example, that holding a portfolio constructed using the optimal weights based on the Random Walk/Linear Regression (RW^{LR}) model yields the same average utility as holding the Forward Premium/Stochastic Volatility (FP^{SV}) optimal portfolio that is subject to monthly expenses Φ, expressed as a fraction of wealth invested in the portfolio. Since the investor would be indifferent between these two strategies, we interpret Φ as the maximum performance fee he will pay to switch from the RW^{LR} to the FP^{SV} strategy. In other words, this utility-based criterion measures how much a mean-variance investor is willing to pay for conditioning on the forward premium under stochastic volatility innovations. The

performance fee will depend on the investor's degree of risk aversion. To estimate the fee, we find the value of Φ that satisfies:

$$\sum_{t=0}^{T-1}\left\{\left(R_{p,t+1}^{*}-\Phi\right)-\frac{\delta}{2\left(1+\delta\right)}\left(R_{p,t+1}^{*}-\Phi\right)^{2}\right\}=\sum_{t=0}^{T-1}\left\{R_{p,t+1}-\frac{\delta}{2\left(1+\delta\right)}R_{p,t+1}^{2}\right\} \qquad (15)$$

where $R_{p,t+1}^{*}$ is the gross portfolio return constructed using the expected return and volatility forecasts from the FP^{SV} model, and $R_{p,t+1}$ is implied by the benchmark RW^{LR} model.

In the context of mean-variance analysis, a commonly used measure of economic value is the Sharpe ratio. However, as suggested by Marquering and Verbeek (2004) and Han (2006), the Sharpe ratio can be misleading because it severely underestimates the performance of dynamic strategies. Specifically, the realized Sharpe ratio is computed using the sample standard deviation of the realized portfolio returns and hence it overestimates the conditional risk an investor faces at each point in time. Furthermore, the Sharpe ratio cannot quantify the exact economic gains of the dynamic strategies over the static random walk strategy in the direct way of the performance fees. Therefore, our economic analysis of short-horizon exchange rate predictability focuses primarily on performance fees, while Sharpe ratios of selected models are reported in the robustness section.[8]

3.4 The Dynamic FX Strategies

In this mean-variance quadratic-utility framework, we design the following global strategy. Consider a US investor who builds a portfolio by allocating his wealth between four bonds: one domestic (US), and three foreign bonds (UK, Germany and Japan). At the beginning of each month, the four bonds yield a riskless return in local currency. Hence the only risk the US investor is exposed to is FX risk. Each month the investor takes two steps. First, he uses each of the 15 models to forecast the one-month ahead conditional mean and volatility of the exchange rate returns. Second, conditional on the forecasts of each model, he dynamically rebalances his portfolio by computing the new optimal weights for the maximum return strategy. This setup is designed to inform us whether using one particular conditional mean and volatility specification affects the performance of a short-horizon allocation strategy in an economically meaningful way. The yields of the riskless bonds are proxied by monthly Eurodeposit rates.

In the context of this maximum return dynamic strategy we compute both the in-sample and the out-of-sample performance fee, Φ, where the out-of-sample period starts in January 1990 and ends in December 2004. Furthermore, we compare the performance fees for the combinations corresponding to the following cases: (i) three sets of target annualized portfolio volatilities $\sigma_{p}^{*}=\{8\%,10\%,12\%\}$; (ii) all pairs of 15 models (for example FP^{SV} vs. RW^{LR}); and (iii) degrees of RRA $\delta=\{2,6\}$. We report the estimates of Φ as annualized fees in basis points.[9]

3.5 Transaction Costs

The impact of transaction costs is an essential consideration in assessing the profitability of trading strategies. This is especially true in our case because the trading strategy based on the random walk benchmark is static (independent of state variables), whereas the remaining empirical models gener-

---

[8]The annualized Sharpe ratios reported in Table 10 are adjusted for the serial correlation in the monthly portfolio returns generated by the dynamic strategies. Specifically, following Lo (2002), we multiply the monthly Sharpe ratios by the adjustment factor $\frac{12}{\sqrt{12+2\sum_{k=1}^{11}(12-k)\rho_{k}}}$, where ρ_{k} is the autocorrelation coefficient of portfolio returns at lag k.

[9]The in-sample period in our economic value results starts in January 1979 due to lack of data for the Japanese Eurocurrency interest rate. In contrast, for the statistical analysis the in-sample period starts in January 1976.

ate dynamic strategies.[10] Furthermore, making an accurate determination of the size of transaction costs is difficult because it involves three factors: (i) the type of investor (e.g. individual vs. institutional investor), (ii) the value of the transaction, and (iii) the nature of the broker (e.g. brokerage firm vs. direct internet trading). This difficulty is reflected in the wide range of estimates used in empirical studies. For example, Marquering and Verbeek (2004) consider three levels of transaction costs, 0.1%, 0.5% and 1%, to represent low, medium and high costs.

Our approach avoids these concerns by calculating the break-even transaction cost, τ^{BE}, that renders investors indifferent between two strategies (e.g. Han, 2006). Hence, we assume that transaction costs equal a fixed proportion (τ) of the value traded in each bond: $\tau |w_t - w_{t-1} \frac{1+r_t}{1+r_{p,t}}|$. In comparing a dynamic strategy with the static (random walk) strategy, an investor who pays transaction costs lower than τ^{BE} will prefer the dynamic strategy. We report τ^{BE} in monthly basis points.[11]

4 Estimation and Forecasting

4.1 Bayesian Markov Chain Monte Carlo Estimation

Stochastic volatility models are generally less popular in empirical applications than GARCH despite their parsimonious structure, intuitive appeal and popularity in theoretical asset pricing. This is primarily due to the numerical difficulty associated with estimating SV models using standard likelihood-based methods because the likelihood function is not available analytically. Bayesian estimation offers a substantial computational advantage over any classical approach because it avoids tackling difficult numerical optimization procedures. In this context, we estimate all three volatility frameworks (LR, GARCH and SV) using similar Bayesian Markov Chain Monte Carlo (MCMC) estimation algorithms. This is a crucial aspect of our econometric analysis because it renders the posterior mean estimates directly comparable across the three volatility structures. It also allows us to use the same model risk diagnostics for all model specifications. Finally, a distinct advantage of Bayesian inference is that it provides the posterior distribution of a regression coefficient conditional on the data, which holds for finite samples and regardless of whether exchange rates (and fundamentals) are (co)integrated (e.g. Sims, 1988).[12]

We estimate the parameters of the SV model using the Bayesian MCMC algorithm of Chib, Nardari and Shephard (2002), which builds on the procedures developed by Kim, Shephard and Chib (1998). The algorithm constructs a Markov chain whose limiting distribution is the target posterior density of the SV parameters. The Markov chain is a Gibbs sampler in which all parameters are drawn sequentially from their full conditional posterior distribution. The Gibbs sampler is iterated 5,000 times and the sampled draws, beyond a burn-in period of 1,000 iterations, are treated as variates from the target posterior distribution. We design a similar Bayesian MCMC algorithm for estimating the GARCH(1,1) parameters, which also draws from the insights of Vrontos, Dellaportas and Politis (2000). The Bayesian Linear Regression algorithm implements a simple MCMC assuming an independent Normal-Gamma prior distribution (for details see Koop, 2003). The MCMC algorithm for each of the three volatility models is summarized in the appendix.

The mean of the MCMC parameter draws is an asymptotically efficient estimator of the posterior

---

[10] The random walk model (RW^{LR}) is the only empirical model that assumes constant mean and variance. Therefore, the in-sample optimal weights for the RW^{LR} trading strategy remain constant over time. However, the out-of-sample optimal weights will vary because every month we re-estimate the drift and variance of the RW^{LR} model.

[11] In contrast to Φ, which is reported in annual basis points, τ^{BE} is reported in monthly basis points because τ^{BE} is a proportional cost paid every month when the portfolio is rebalanced.

[12] This is not the case in classical inference, where the small samples typically employed in the study of exchange rate predictability combined with the assumption that exchange rates and fundamentals are cointegrated can have a critical impact in overstating predictability (e.g. Berkowitz and Giorgianni, 2001).

mean of θ (see Geweke, 1989). The Numerical Standard Error (NSE) is the square root of the asymptotic variance of the MCMC estimator:

$$NSE = \sqrt{\frac{1}{I}\left\{\widehat{\psi}_0 + 2\sum_{j=1}^{B_I} K\left(z\right)\widehat{\psi}_j\right\}} \tag{16}$$

where $I = 5,000$ is the number of iterations (beyond the initial burn-in of $1,000$ iterations), $j = 1,...,B_I = 500$ lags is the set bandwidth, $z = \frac{j}{B_I}$, and $\widehat{\psi}_j$ is the sample autocovariance of the MCMC draws for each estimated parameter cut according to the Parzen kernel $K\left(z\right)$.

The likelihood function of the SV models is not available analytically, and hence must be simulated. The log-likelihood function is evaluated under the predictive density as:

$$\log\widehat{L} = \sum_{t=1}^{T}\log\widehat{f}\left(\Delta s_t \mid \mathcal{F}_{t-1}, \theta\right) = \sum_{t=1}^{T}\log\widehat{f}_t\left(\Delta s_t \mid h_t, \theta\right) \tag{17}$$

where θ is taken as the posterior mean estimate from the MCMC simulations. The key to this calculation is simulating the one-step ahead predictive log-variance $h_t \mid \mathcal{F}_{t-1}, \theta$, which is a non-trivial task as it is sampled using the particle filter of Pitt and Shephard (1999). The particle filter is summarized in the appendix.

4.2   Model Risk and Posterior Probability

Model risk arises from the uncertainty over selecting a model specification. Consistent with our Bayesian approach, a natural statistical criterion for resolving this uncertainty is the posterior probability of each model. Hence, we rank the competing models using the posterior probability, which has three important advantages relative to the log-likelihood: (i) it is based on the marginal likelihood and therefore accounts for parameter uncertainty, (ii) it imposes a penalty for lack of parsimony (higher dimension), and (iii) it forms the basis of the Bayesian Model Averaging strategy discussed below. Ranking the models using the highest posterior probability is equivalent to choosing the best model in terms of density forecasts and is a robust model selection criterion in the presence of misspecification and non-nested models (e.g. Fernandez-Villaverde and Rubio-Ramirez, 2004).

Consider a set of N models $M_1, ..., M_N$. We form a prior belief $\pi\left(M_i\right)$ on the probability that the ith model is the true model, observe the FX returns data Δs, and then update our belief that the ith model is true by computing the posterior probability of each model defined as follows:

$$p\left(M_i \mid \Delta s\right) = \frac{p\left(\Delta s \mid M_i\right)\pi\left(M_i\right)}{\sum_{j=1}^{N} p\left(\Delta s \mid M_j\right)\pi\left(M_j\right)} \tag{18}$$

where $p\left(\Delta s \mid M_i\right)$ is the marginal likelihood of the ith model defined as follows:

$$p\left(\Delta s \mid M_i\right) = \int_{\theta} p\left(\Delta s, \theta \mid M_i\right)d\theta = \int_{\theta} p\left(\Delta s \mid \theta, M_i\right)\pi\left(\theta \mid M_i\right)d\theta \tag{19}$$

In Equation 18 above we set our prior belief to be that all models are equally likely, i.e. $\pi\left(M_i\right) = \frac{1}{N}$.

Note that the marginal likelihood is an averaged (not a maximized) likelihood. This implies that the posterior probability is an automatic "Occam's Razor" in that it integrates out parameter uncertainty.[13] Furthermore, the marginal likelihood is simply the normalizing constant of the posterior

---

[13] *Occam's Razor* is the principle of parsimony, which states that among two competing theories that make exactly

density and (suppressing the model index for simplicity) it can be written as:

$$p\left(\Delta s\right) = \frac{f\left(\Delta s \mid \theta\right) \pi\left(\theta\right)}{\pi\left(\theta \mid \Delta s\right)} \tag{20}$$

where $f\left(\Delta s \mid \theta\right)$ is the likelihood, $\pi\left(\theta\right)$ the prior density of the parameter vector θ, $\pi\left(\theta \mid \Delta s\right)$ the posterior density, and θ is evaluated at the posterior mean. Since θ is drawn in the context of MCMC sampling, the posterior density $\pi\left(\theta \mid \Delta s\right)$ is computed using the technique of reduced conditional MCMC runs of Chib (1995) and Chib and Jeliazkov (2001).

4.3 Combined Forecasts

Assessing the predictive ability of empirical exchange rate models primarily involves a pairwise comparison of the competing models. However, given that we do not know which one of the models is true, it is important that we assess the performance of combined forecasts proposed by the seminal work of Bates and Granger (1969). Specifically, we design two strategies based on a combination of forecasts for both the conditional mean and volatility of exchange rate returns: the Bayesian Model Average (BMA) strategy and the Bayesian Winner (BW) strategy.[14]

We assess the economic value of combined forecasts by treating the BMA and BW strategies in the same way as any of the 15 individual empirical models. For instance, we compute the performance fee, Φ, for the BMA one-month ahead forecasts and compare it to the random walk benchmark. We focus on two distinct universes of models: the restricted universe of the five SV models (because the five conditional mean specifications with SV innovations have the highest marginal likelihood), and the unrestricted universe of all 15 empirical exchange rate models.

4.3.1 The BMA Strategy

In the context of our Bayesian approach, it is natural to implement the BMA method originally discussed in Leamer (1978) and surveyed in Hoeting, Madigan, Raftery and Volinsky (1999). The BMA strategy accounts directly for uncertainty in model selection, and is in fact easy to implement once we have the output from the MCMC simulations. Define $f_{i,t}$ as the forecast density of each of the N competing models at time t. Then, the BMA forecast density is given by:

$$f_t^{BMA} = \sum_{i=1}^{N} p_t\left(M_i \mid \Delta s_t\right) f_{i,t} \tag{21}$$

where $p_t\left(M_i \mid \Delta s_t\right)$ is the posterior probability of model M_i given the data Δs_t.

It is important to note that: (i) the BMA weights vary not only across models but also across time periods as does the marginal likelihood of each model, and (ii) we evaluate the BMA strategy ex-ante. We do this by lagging the posterior probability of each model for the following reason. Suppose that we need to compute the period t BMA forecasts of the conditional mean and volatility for the four bonds we include in the portfolio. Knowing the mean and volatility forecasts implied by each model for the three exchange rates is not sufficient. We also need the realized data point Δs_t in order to evaluate the predictive density $f_t\left(\Delta s_t \mid \mathcal{F}_{t-1}, \theta\right)$. Since the realized data point Δs_t is only

---

the same prediction, the simpler one is best.

[14] See Diebold and Pauly (1990), Diebold (1998, 2004) and Timmermann (2006) for a review of forecast combinations. A previous version of the paper also considers a deterministic model average (DMA) method, which involves taking an equally weighted average of the conditional mean and volatility forecasts from a given universe of available models; we find that the BMA and BW combination methods outperform the DMA method (results available upon request).

observed ex post, the only way to form the BMA weights ex ante is to lag the predictive density and thus use $f_{t-1}\left(\Delta s_{t-1} \mid \mathcal{F}_{t-2}, \theta\right)$.

4.3.2 The BW Strategy

Under the BW strategy, in each time period we select the set of one-step ahead conditional mean and volatility from the empirical model that has the highest marginal likelihood up to that period. In other words, the BW strategy only uses the forecasts of the "winner" model in terms of marginal predictive density, and hence discards the forecasts of the rest of the models. Clearly, there is no model averaging in the BW strategy. Similar to the BMA, the BW strategy is evaluated ex ante using the lagged marginal likelihood.

5 Empirical Results

5.1 FX Data and Descriptive Statistics

The data sample consists of 348 monthly observations ranging from January 1976 to December 2004, and focuses on three exchange rates relative to the US dollar: the UK pound sterling (USD/GBP), Deutsche mark/euro (USD/DEM-EURO), and Japanese yen (USD/JPY). The spot and one-month forward exchange rates are taken from *Datastream* for the period of January 1985 onwards, whereas for the period ranging from January 1976 to December 1984 they are taken from Hai, Mark and Wu (1997). After the introduction of the euro in January 1999, we use the euro exchange rate to replace the Deutsche mark rate.

Data on money supply and income are from the International Monetary Fund's *International Financial Statistics* database. Specifically, we define the money supply as the sum of money (line code 34) and quasi-money (line code 35) for Germany and Japan, whereas for the UK we use $M0$ (line code 19). Since German exchange rate data are only available until December 1998, we use the money and quasi-money data of the Euro Area for the remaining period (January 1999 to December 2004). The US data is obtained from the aggregate $M2$ of the Board of Governors of the Federal Reserve System. Furthermore, we use the monthly industrial production index (line code 66) as a proxy for national income rather than the gross domestic product (GDP), because the latter is available only at the quarterly frequency.[15] We deseasonalize the money and industrial production indices following the procedure of Gomez and Maravall (2000). Note that we ignore the complication arising from the fact that the data we use on monetary fundamentals may not be available in real time and may not suffer from the measurement errors that characterize real-time macroeconomic data (Faust, Rogers and Wright, 2003). This issue will not affect our main findings on the predictive ability of the forward premium and stochastic volatility.

We take logarithmic transformations of the raw data to yield time series for s_t, f_t, m_t, m_t^*, y_t and y_t^*. The monetary fundamentals series, z_t, is constructed as in Equation 2; s_t is taken as the natural logarithm of the domestic price of foreign currency, the US being the domestic country; f_t is the natural logarithm of the US dollar price of a one-month forward contract issued at time t for delivery of one unit of foreign currency at time $t + 1$. Finally, in our economic evaluation of the set of candidate exchange rate models, the proxy for the riskless domestic and foreign bonds is the end-of-month Euromarket interest rate with one month maturity, obtained from *Datastream*.[16]

---

[15] For all countries, the correlation coefficient between the quarterly industrial production index and GDP over our sample period is higher than 0.95.

[16] We use the Eurocurrency deposit rate as a proxy for the riskless rate because these deposits are comparable across countries in all respects (such as credit risk and maturity) except for currency of denomination; see Levich (1985).

Table 1 reports the descriptive statistics for the monthly percent FX returns, Δs_t, the three monetary fundamentals predictors, MF_1, MF_2 and MF_3, also expressed in percent, and the percent forward premium, $f_t - s_t$. For our sample period, the sample means of the FX returns are -0.012% for USD/GBP, 0.165% for USD/DEM-EURO and 0.309% for USD/JPY. The FX return standard deviations are similar across the three exchange rates at about 3% per month. Finally, the exchange rate return sample autocorrelations are approximately 0.10 but decay rapidly.

The three specifications of monetary fundamentals predictors display high volatility and persistence. For instance, the standard deviation of MF_1 is about 20% for the UK, 30% for Germany and 40% for Japan. However, the standard deviation of MF_3 (which is corrected for both the intercept and the time trend component) is approximately half the value of the canonical monetary fundamentals MF_1. The three monetary fundamentals predictors exhibit low skewness, low excess kurtosis and high serial correlation. Finally, the average forward premium is negative for the UK, but positive for Germany and Japan. The standard deviation of $f_t - s_t$ is low across all exchange rates (in fact, about 100 times smaller than MF_1), but the forward premium exhibits high kurtosis and its sample autocorrelation is high and decreasing slowly.

5.2 Estimation of Exchange Rate Models

We begin our statistical and economic evaluation of short-horizon exchange rate predictability by performing Bayesian estimation of the parameters of our 15 candidate models: the five conditional mean specifications (RW, MF_1, MF_2, MF_3 and FP) under the three volatility frameworks (LR, $GARCH$ and SV). The posterior mean estimates for the model parameters are presented in Tables 2, 3 and 4. We particularly focus on the size, sign and statistical significance of the β estimate because it captures the effect of either monetary fundamentals or the forward premium in the conditional mean of exchange rate returns. In our Bayesian MCMC framework we assess statistical significance using two diagnostics. First, we report the highest posterior density (HPD) region for each parameter estimate. For example, the 95% HPD region is the shortest interval that contains 95% of the posterior distribution. We check whether the 90%, 95% and 99% HPD regions contain zero, which is equivalent to two-sided hypothesis testing at the 10%, 5% and 1% level respectively. Second, we compute the Numerical Standard Error (NSE) as defined in Equation 16.

Tables 2 through 4 illustrate that for the three monetary fundamentals specifications (MF_1, MF_2 and MF_3) the in-sample β estimate tends to be a low positive number, which increases in size as we move from MF_1 to MF_3. This suggests that when s_t is below (above) its fundamental value z_t, it is expected to slowly rise (decrease) over time. In contrast, the in-sample β estimate for the FP model has a large negative value. The tables also report the estimates of the conditional variance parameters. For the Linear Regression model, the monthly variance of FX returns remains largely unchanged across the five conditional mean specifications and is around 10 (i.e. $\upsilon \approx 3\%$) for all three currencies. For the $GARCH(1,1)$ models, the conditional monthly variance is highly persistent since the sum $\gamma_1 + \gamma_2$ revolves around 0.96 for all specifications. The SV models exhibit (i) high persistence (ϕ) in the conditional monthly log-variance, and (ii) a sizeable stochastic component in the conditional monthly log-variance. Finally, all parameters in both the conditional mean and volatility exhibit very low NSE values and therefore a high degree of statistical significance.

5.3 Evaluating Forecasts Using Statistical Criteria

We assess the statistical evidence on short-horizon exchange rate predictability by ranking our set of 15 candidate models according to their log-likelihood and posterior probability. The conditional performance of the models is evaluated in-sample as well as out-of-sample. The in-sample period for

the three monthly exchange rates covers 29 years ranging from January 1976 to December 2004. The out-of-sample exercise involves two steps: (i) initial parameter estimation for the 14-year period of January 1976 to December 1989, and (ii) sequential monthly updating of the parameter estimates for the out-of-sample 15-year period of January 1990 to December 2004. In other words, the forecasts at any given month are constructed according to a recursive procedure that is conditional only upon information up to the date of the forecast. The model is then successively re-estimated as the date on which forecasts are conditioned moves through the data set. Hence the design of the out-of-sample exercise is computationally intensive.

Our analysis of the statistical evidence begins with Table 5, which presents the log-likelihood values and shows that across volatility models, the SV model always has higher log-likelihood than both LR and $GARCH$. This result is robust as it holds for all currencies both in-sample and out-of-sample. Similarly, the $GARCH(1,1)$ model always beats the constant variance LR models in terms of log-likelihood. Across conditional mean specifications, the RW model is always worse in-sample. Finally, the out-of-sample log-likelihood values lead to the following conclusions: FP is the best model for the yen, but the RW model is best for the pound sterling and the Deutsche mark/euro.

In Table 6 we rank the in-sample and out-of-sample performance of our set of candidate models according to their posterior probability. The key input to this statistical criterion is the calculation of the marginal likelihood. Therefore, Table 6 gives us a distinct statistical perspective on performance because the marginal likelihood is computed in a way that integrates out parameter uncertainty and imposes a penalty for lack of parsimony (higher dimension). The results in Table 6 indicate two clear patterns in ranking the models. The first pattern confirms one of our most robust results: the best models for all three currencies both in-sample and out-of-sample have SV innovations. The second pattern is slightly different from the log-likelihood findings: for all three exchange rates, both in-sample and out-of-sample, the best model is FP^{SV}, the second best is RW^{SV}, and third best is one of the three MF^{SV} specifications. The single exception is the pound sterling for which RW^{SV} is the best out-of-sample model. Hence, in contrast to the likelihood evidence, the MF specifications lose to RW even in-sample. In other words, the penalty the posterior probability imposes on the three monetary fundamentals models for lack of parsimony offsets their log-likelihood advantage.

5.4 Evaluating Forecasts Using Economic Criteria

We assess the economic value of short-horizon exchange rate predictability by analyzing the performance of the dynamically rebalanced portfolios constructed using our set of 15 candidate models. Our analysis focuses on the performance fee, Φ, a US investor is willing to pay for switching from one FX strategy to another. The fees are reported in Table 7, which displays the economic value of each mean and volatility specification relative to the benchmark random walk model with constant variance (RW^{LR}). We present the fees for the degrees of RRA $\delta = 2$ and $\delta = 6$.

Panel A of Table 7 presents the in-sample performance fees and demonstrates that the three monetary fundamentals specifications generally have no economic value as indicated by the negative Φ values. In contrast, the forward premium model (FP) exhibits high economic value, especially under stochastic volatility. For example, at the target portfolio volatility of $\sigma_p^* = 10\%$ and for $\delta = 2$, a US investor is willing to pay a substantial 248 annual basis points (bps) for switching from the RW^{LR} model to FP^{SV}. Consistent with our statistical evidence, for all conditional mean specifications there tends to be high economic value associated with stochastic volatility. However, contrary to our statistical evidence, the performance of the $GARCH(1,1)$ model is surprisingly poor relative to the constant variance Linear Regression model. For $\sigma_p^* = 10\%$ and $\delta = 2$, the in-sample fee for switching from RW^{LR} to RW^{GARCH} is -24 bps, whereas the fee for switching from RW^{LR} to

RW^{SV} is 42 bps.[17] Finally, as investors become less risk averse, the fees tend to increase in absolute value, strengthening the evidence against RW^{LR} and in favor of the FP^{SV} specification.

The out-of-sample performance fees are displayed in Panel B of Table 7 and suggest that even out-of-sample there is still high economic value in both the forward premium and stochastic volatility. This is a new and important result, which adds to the existing literature that is anchored around the seminal contribution of Meese and Rogoff (1983). Specifically, at $\sigma_p^* = 10\%$ and $\delta = 2$, the annual performance fees for switching from RW^{LR} to another model are 127 bps for RW^{SV} and 266 bps for FP^{SV}. We can therefore conclude that there is substantial economic value both in-sample and out-of-sample against the naive random walk model and in favour of conditioning on the forward premium with stochastic volatility. This finding is in fact consistent with the large profits made by financial institutions that engage in sophisticated multi-currency forward bias strategies. For example, Galati and Melvin (2004) show that simple carry trades aiming at exploiting the forward bias constitute a significant source of the surge in FX trading in recent years.

In addition to the results associated with individual models, even stronger economic evidence is found for the combined forecasts reported in Table 8, which compares BMA and BW to the RW^{LR} benchmark for two cases: (i) the restricted universe of the five SV models (because the SV models generally perform better), and (ii) the unrestricted universe of all 15 models. A purely agnostic approach to forecast combination would use the full set of 15 models (case ii). The results in Table 8 provide robust evidence against the naive random walk model as all performance fees based on the BMA and BW are positive and high, both in-sample and out-of-sample. For example, when selecting among the SV models and setting $\sigma_p^* = 10\%$ and $\delta = 2$, the annual in-sample performance fee for switching away from the benchmark RW^{LR} is 255 bps for BMA and 235 bps for BW. The out-of-sample fees are even higher at 317 bps for BMA and 340 bps for BW. In short, therefore, there is clear in-sample and out-of-sample economic evidence on the superiority of combined forecasts relative to the naive random walk benchmark.

In conclusion, Figure 1 offers a visual description of the time variation in the weights investing in the three risky assets: the UK, German and Japanese bonds. The figure displays the weights for four cases: the benchmark RW^{LR} model, the best performing individual model FP^{SV}, the BMA and the BW combined forecast strategies. As expected, the weights are very smooth over time for RW^{LR}, and remain reasonably smooth for the FP^{SV} model and the two combined forecast strategies.[18]

5.5 Transaction Costs

If transaction costs are sufficiently high, the period-by-period fluctuations in the dynamic weights of an optimal strategy will render the strategy too costly to implement relative to the static random walk model. We address this concern by computing the break-even transaction cost, τ^{BE}, as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. In comparing a dynamic strategy with the static random walk strategy, an investor who pays a transaction cost lower than τ^{BE} will prefer the dynamic strategy. The τ^{BE} values are expressed in monthly basis points and are reported only when Φ is positive.

The in-sample break-even transaction costs are reported in Panel A of Table 7, which demon-

---

[17]At first sight, the poor performance of the GARCH model in terms of economic value appears rather surprising. For instance, Fleming and Kirby (2003) find that SV models only marginally outperform GARCH models. However, there is no study to date which assesses the economic value of GARCH and SV models, especially when applied to exchange rates. Furthermore, the negative in-sample and out-of-sample performance fees of RW^{GARCH} are not far from zero.

[18]However, the dynamic weights appear to be more volatile in the beginning of the sample before they stabilize. We believe that the initial instability in the weights is due to the high exchange rate volatility around the 1992 crisis of the Exchange Rate Mechanism that forced the UK to abandon the target zone system.

strates that for the forward premium and stochastic volatility the values of τ^{BE} are positive and high; they tend to be higher than 100 bps and can be as high as 556 bps. For instance, at $\sigma_p^* = 10\%$ and $\delta = 2$, a US investor will switch back to the RW^{LR} model if he is subject to a proportional transaction cost of at least 120 bps for FP^{LR}, 101 bps for FP^{GARCH}, 132 bps for FP^{SV}, and 471 bps for RW^{SV}. In other words, at the reasonably high transaction cost of 50 bps (e.g. Marquering and Verbeek, 2004), there is still significant in-sample economic value in empirical models which condition on the forward premium, especially under stochastic volatility.

Determining the out-of-sample robustness to transaction costs is one of the most important considerations in assessing the forecasting performance of empirical exchange rate models. Panel B of Table 7 shows that conditioning on the forward premium and stochastic volatility leads to reasonably high τ^{BE} values. Specifically, at $\sigma_p^* = 10\%$ and $\delta = 2$, the break-even transaction cost which would eliminate the performance fee of 266 bps of the FP^{SV} model relative to the RW^{LR} benchmark is 90 bps. Furthermore, the τ^{BE} for RW^{SV} versus RW^{LR} is a very large 321 bps.

The evidence on the τ^{BE} of combined forecasts displayed in Table 8 is even stronger. Compared to the benchmark RW^{LR}, a combined forecast of all 15 models exhibits an in-sample τ^{BE} of 141 bps for BMA and 114 bps for BW. Panel B of Table 8 shows that the out-of-sample τ^{BE} values for combined forecasts are generally as high as in sample. In short, as the τ^{BE} values are generally positive and reasonably high, we conclude that the in-sample and out-of-sample economic value we have reported is robust to reasonably high transaction costs for empirical exchange rate models conditioning on the forward premium, for models with SV innovations, and for combined forecasts.

5.6 Summary of Results

The statistical and economic evidence on short-horizon exchange rate predictability supports the following four results: (i) the forward premium model unequivocally beats the random walk; (ii) conditioning on monetary fundamentals has no economic value; (iii) the stochastic volatility process always leads to superior portfolio performance; and (iv) the combined forecasts consistently outperform the constant variance random walk benchmark. All these results hold both in-sample and out-of-sample and are robust to reasonably high transaction costs.

6 Robustness and Extensions

This section discusses directions in which one can possibly extend the analysis of the paper. First, we perform an additional robustness test by evaluating the out-of-sample performance of the empirical models in three 5-year subsamples. Recall that the full sample period at our disposal covers 29 years ranging from January 1976 to December 2004. We use data from January 1976 to December 1989 for in-sample estimation, whereas the out-of-sample period contains 15 years ranging from January 1990 to December 2004. The out-of-sample results we report in Tables 5 through 8 are for the entire 15-year out-of-sample period. In addition, Panel A of Table 9 presents the performance fees for selected models and for three subsamples: 1990-1994, 1995-1999 and 2000-2004. We find that the economic value of conditioning on the forward premium and stochastic volatility is positive in all periods but is substantially higher in the last two subsamples. This is consistent with the well-known fact in the literature that the forward bias is small in the early 1990s (e.g. Flood and Rose, 2002).[19] For all models, the best subsample period is 1995-1999. Furthermore, it is important to note that

---

[19]In a separate experiment we start the out-of-sample exercise in 1985 and find significant economic value in the forward premium and stochastic volatility for the 1985-1989 period. However, starting the out-of-sample period in 1985 leaves too few in-sample observations for initial parameter estimation. Therefore, the tables present the out-of-sample results for the period starting in 1990.

the combined forecast strategies substantially outperform the random walk benchmark in all three subsamples and display similar performance fees to FP^{SV} for the last two subsamples. However, for the first subsample when the forward bias is small, the BMA and BW strategies significantly outperform FP^{SV} by optimally using predictive information from the entire universe of models, including monetary fundamentals.

Second, our analysis of the conditional variance of exchange rate returns includes the $GARCH(1,1)$ specification because this is the benchmark model in the seminal study of West, Edison and Cho (1993). As a further robustness check, we examine the out-of-sample performance of the $tGARCH(1,1)$ model of Bollerslev (1987) in order to determine whether departing from the assumption of conditional normality improves the performance of the GARCH model.[20] The results from this exercise (a subset of which is reported in Panel B of Table 9) reveal that using a Student-t distribution leads to substantial performance gains in the GARCH framework.[21] In particular, we find that the out-of-sample performance fees of the $tGARCH(1,1)$ model are much higher than for $GARCH(1,1)$, especially for the forward premium and random walk conditional mean specifications. For instance, setting $\sigma_p^* = 10\%$ and $\delta = 2$ and comparing the results in Tables 7 and 9 indicates that the out-of-sample fees for switching from the RW^{LR} model to the forward premium models are as follows: 76 bps for FP^{LR}, 70 bps for FP^{GARCH}, 140 bps for FP^{tGARCH}, and 266 bps for FP^{SV}. Similarly, when switching from the random walk with constant variance, RW^{LR}, to a random walk with time-varying volatility the fees are: -32 bps for RW^{GARCH}, 28 bps for RW^{tGARCH}, and 127 bps for RW^{SV}. Therefore, we can conclude that in terms of economic value the $tGARCH$ model performs better than $GARCH$, although the SV model outperforms both normal and Student-t GARCH specifications. Hence, our main conclusions remain qualitatively the same.

Third, Table 10 presents the in-sample and out-of-sample annualized Sharpe ratios for selected models. The Sharpe ratio values are generally in agreement with the performance fees and hence confirm our conclusions. Specifically, FP^{SV}, BMA and BW consistently outperform RW^{LR} both in-sample and out-of-sample. For example, at $\sigma_p^* = 10\%$, the out-of-sample Sharpe ratios are: 0.76 for RW^{LR}, 0.98 for FP^{SV}, 1.06 for BMA, and 1.12 for BW.

Fourth, this paper explores the predictability in exchange rates by focusing on the frequency and horizon of one month. On the one hand, adopting the monthly frequency is a natural choice because this is the highest frequency at which monetary fundamentals are observed. On the other hand, our motivation for investigating predictability at the one-month horizon is founded on the prevailing view in this literature that exchange rates are not predictable at short horizons. It is clear, therefore, that one possible direction in extending the analysis of this paper is to study the predictability of the forward premium, stochastic volatility and combined forecasts for higher frequencies and longer horizons. We leave this for future research.

Finally, we study short-horizon exchange rate predictability by estimating a set of univariate conditional mean and volatility models. However, in assessing the economic value of exchange rate predictability we build multivariate dynamic asset allocation strategies. Specifically, the optimal weights of the dynamically rebalanced portfolios are computed using the conditional mean forecasts, the conditional volatility forecasts and the dynamic covariances implied by the constant conditional correlation (CCC) model of Bollerslev (1990). In the CCC model, the dynamics of covariances are driven by the time-variation in the conditional volatilities. By design, therefore, the advantage of this setting is that the optimal weights will vary across models only to the extent that forecasts of the con-

---

[20]In estimating the $tGARCH$ model, we implement an algorithm similar to the $GARCH$ case as described in Appendix A.3, with an additional Metropolis-Hastings step for sampling the degrees of freedom parameter ν.

[21]Note that the degrees of freedom parameter estimate revolves around $\nu = 10$ for the UK pound, $\nu = 25$ for the Deutsche mark/euro, and $\nu = 7$ for the Japanese yen (not reported). This indicates that the unconditional distribution of exchange rate returns is not normal, especially for the UK pound and the Japanese yen.

ditional mean and volatility will vary, which is precisely what the empirical models provide. Indeed, introducing multivariate stochastic volatility models for capturing the dynamic heteroskedasticity of the covariances of exchange rate returns remains an important extension to this line of research. Multivariate stochastic volatility models are high dimensional and their estimation is computationally challenging (e.g. Chib, Nardari and Shephard, 2006). Additionally the dynamic conditional correlation (DCC) model of Engle (2002) has yet to be examined in a Bayesian framework. We will revisit this issue in future research.

7   Conclusion

This paper draws from three separate, yet related strands of international finance literature. A large body of empirical research finds that models which condition on monetary fundamentals cannot outperform the naive random walk model in out-of-sample forecasting of exchange rates. Despite the increasing sophistication of the econometric techniques implemented and the improving quality of the data sets utilized, evidence of exchange rate predictability remains elusive. A second and related research strand indicates that the rejection of the risk-neutral FX efficient market hypothesis implies that exchange rate movements can be predicted using information contained in forward premia. Finally, financial economists agree that exchange rate volatility is predictable by specifying either GARCH or stochastic volatility innovations.

Prior research in this area has largely relied on standard statistical measures of forecast accuracy. In this paper, we complement this approach in two critical aspects. First, in assessing the predictive performance of the set of empirical exchange rate models, we implement a Bayesian methodology which explicitly accounts for parameter and model uncertainty. Second, we provide a comprehensive economic evaluation of the models in the context of dynamic asset allocation strategies. In doing so, our study contributes to the growing empirical literature on exchange rate predictability in the following manner. We assess the economic value of exchange rate forecasts derived from empirical models which condition on information contained in either monetary fundamentals or forward premia. This is done in a framework that allows for time-varying volatility. The empirical exchange rate models are set against the naive random walk benchmark. Finally, we evaluate the performance of combined forecasts based on Bayesian Model Averaging.

Our results provide robust evidence against the random walk (no predictability) benchmark, and therefore our empirical findings reinforce the notion that exchange rates are predictable. Specifically, we find that the predictive ability of the forward premium has substantial economic value in a dynamic portfolio allocation context and that stochastic volatility significantly outperforms the constant variance and GARCH(1,1) models irrespective of the conditional mean specification. Combined forecasts formed using Bayesian Model Averaging also substantially outperform the random walk. These results are robust to reasonably high transaction costs and hold for all currencies both in-sample and out-of-sample. In short, these findings suggest that the random walk hypothesis as applied to exchange rates might have been overstated, while at the same time they justify the widespread use of forward bias and volatility timing strategies in the practice of currency management.

A   APPENDIX: Bayesian MCMC Estimation

We perform Bayesian MCMC estimation of the parameters of the empirical exchange rate models by constructing a Markov chain whose limiting distribution is the target posterior density. This Markov chain is a Gibbs sampler in which all parameters are drawn sequentially from their full conditional posterior distribution. The chain is then iterated and the sampled draws, beyond a burn-in period, are treated as variates from the target posterior distribution.

A.1   The Linear Regression Algorithm

In the Bayesian Linear Regression (LR) model, we need to estimate $\theta = \{\theta_1, \theta_2\}$, where $\theta_1 = \{\alpha, \beta\}$ is the set of the conditional mean parameters, and $\theta_2 = \{v^{-2}\}$ is the constant precision defined as the inverse of the variance. We define the following priors: for $\theta_1 = \{\alpha, \beta\}$ we assume a Normal prior $N(\underline{\theta}_1, \ \underline{V})$, where $\underline{\theta}_1 = 0_2$ and $\underline{V} = I_2$; for $\theta_2 = \{v^{-2}\}$ we assume a prior $Gamma\left(\frac{\nu}{2}, \frac{2\underline{s}^{-2}}{\underline{\nu}}\right)$ with mean $\underline{s}^{-2} = 1$, and degrees of freedom $\underline{\nu} = 2$. The posterior distributions are summarized in the following simple Gibbs algorithm (for more details see Koop, 2003):

1. Initialize θ_2.

2. Sample θ_1 from $\theta_1 \mid \Delta s, \theta_2 \sim N\left(\overline{\theta}_1, \overline{V}\right)$, where $\overline{V} = \left(\underline{V}^{-1} + \theta_2 X'X\right)^{-1}$ and $\overline{\theta}_1 = \overline{V}\left(\underline{V}^{-1}\underline{\theta}_1 + \theta_2 X'\Delta s\right)$.

3. Sample θ_2 from $\theta_2 \mid \Delta s, \theta_1 \sim Gamma\left(\frac{\overline{\nu}}{2}, \frac{2\overline{s}^{-2}}{\overline{\nu}}\right)$, where $\overline{\nu} = T + \underline{\nu}$ and $\overline{s}^2 = \frac{(\Delta s - X\theta_1)'(\Delta s - X\theta_1) + \underline{\nu}\underline{s}^2}{\overline{\nu}}$.

4. Go to step 2 and iterate $100,000$ times beyond a burn-in of $20,000$ iterations.

A.2   The GARCH(1,1) Algorithm

In the Bayesian GARCH(1,1) model, we need to estimate $\theta = \{\theta_1, \theta_2\}$, where $\theta_1 = \{\alpha, \beta\}$ is the set of the conditional mean parameters, and $\theta_2 = \{\omega, \gamma_1, \gamma_2\}$ are the conditional variance parameters. We ensure that the conditional variance is covariance stationary by specifying ω as a lognormal prior: $\omega \sim LogN\left(\underline{w}, \underline{W}\right)$, with $\underline{w} = -1$ and $\underline{W} = 2$. The prior specification is completed by assuming $\gamma_1 \sim Beta\left(\underline{g}_1, \underline{G}_1\right)$ and $\gamma_2 \sim Beta\left(\underline{g}_2, \underline{G}_2\right)$, where $\underline{g}_1 = 40$, $\underline{G}_1 = 5$, $\underline{g}_2 = 2$, and $\underline{G}_2 = 40$. These hyperparameters imply a mean of 0.89 and 0.05 for γ_1 and γ_2, respectively. The algorithm is summarized below:

1. Initialize θ_1 and transform the data into $\Delta s_t^* = (\Delta s_t - \alpha - \beta x_{t-1})$.

2. Sample the variance parameters θ_2 from their full conditional posterior density: $\theta_2 \mid \Delta s^*, \theta_1$. This posterior density is not available analytically. We compute the log-likelihood of the transformed data Δs_t^* as function of θ_2 (conditional on θ_1) and then we optimize the conditional log-posterior. We generate a proposal from a t-distribution $t(m, V, \xi)$, where m is the mode, V is the inverse of the negative Hessian, and ξ a tuning parameter. The proposal is then accepted according to the independence chain Metropolis-Hasting algorithm (e.g. Chib and Greenberg, 1995).

3. Sample all the conditional mean coefficients $\theta_1 \mid \Delta s, \theta_2$ using a precision-weighted average of a set of normal priors and the normal likelihood conditional on θ_2.

4. Update the data $\Delta s_t^* = (\Delta s_t - \alpha - \beta x_{t-1})$.

5. Go to step 2 and iterate $5,000$ times beyond a burn-in of $1,000$ iterations.

A.3 The Stochastic Volatility Algorithm

In the Bayesian SV model, we need to estimate $\theta = \{\theta_1, \theta_2\}$, where $\theta_1 = \{\alpha, \beta\}$ is the set of the conditional mean parameters, and $\theta_2 = \{\mu, \phi, \sigma^2\}$ are the conditional log-variance parameters. The prior for μ is $N(\underline{m}, \underline{M})$ with $\underline{m} = 1$ and $\underline{M} = 25$. Following Kim, Shephard and Chib (1998), we formulate the prior for ϕ in terms of $\phi = 2\phi^* - 1$, where ϕ^* is distributed as $Beta(\underline{f}, \underline{F})$. This implies that the prior on $\phi \in (-1, 1)$ is $p(\phi) = \kappa\{0.5(1 + \phi)\}^{\underline{f}-1}\{0.5(1 - \phi)\}^{\underline{F}-1}$, $\underline{f}, \underline{F} > 0.5$, where $\kappa = 0.5\frac{\Gamma(\underline{f}+\underline{F})}{\Gamma(\underline{f})+\Gamma(\underline{F})}$. Specifying $\underline{f} = 20$ and $\underline{F} = 1.5$ yields a mean of 0.86 with a variance of 0.01. For σ, the prior is inverse gamma $IG(\underline{s}, \underline{S})$ with $\underline{s} = 3$ and $\underline{S} = 2.5$ so that the distribution has a mean of 0.20 with a variance of 0.006. The parameters of the SV model are estimated using the Bayesian MCMC algorithm of Chib, Nardari and Shephard (2002), which builds on the procedures developed by Kim, Shephard and Chib (1998), and is summarized below:

1. Initialize θ, mx, and transform the data into $\Delta s_t^* = \ln\left((\Delta s_t - \alpha - \beta x_{t-1})^2 + c\right)$, $c = 0.001$ to put the model in state-space form. The "offset" constant c eliminates the inlier problem.

2. Sample the log-variance parameters θ_2 from their full conditional posterior density: $\theta_2 \mid \Delta s^*, mx$. This posterior density is not available analytically. We use the Kalman filter to compute the log-likelihood of the transformed data Δs_t^* as a function of θ_2 (conditional on mx_t) and then optimize the conditional log-posterior. We generate a proposal from a t-distribution $t(m, V, \xi)$, where m is the mode, V is the inverse of the negative Hessian, and ξ a tuning parameter. The proposal is then accepted according to the independence chain Metropolis-Hastings algorithm (e.g. Chib and Greenberg, 1995).

3. Sample the log-variance vector $\{h_t\}$ in one block from the posterior distribution: $h \mid \Delta s^*, mx, \theta_2$. This step uses the de Jong and Shephard (1995) simulation smoother, which is an algorithm designed for efficient sampling of the state vector in a state-space model. See also Carter and Kohn (1994).

4. Sample all the conditional mean coefficients θ_1 from $\theta_1 \mid \Delta s, h$ using a precision-weighted average of a set of normal priors and the normal likelihood conditional on h. Then update the transformed data $\Delta s_t^* = \ln\left((\Delta s_t - \alpha - \beta x_{t-1})^2 + c\right)$, $c = 0.001$.

5. Finally, sample the mixture indicator variable $mx \mid \Delta s^*, h, \theta$ directly from its posterior:

$$\Pr(mx_t \mid \Delta s_t^*, h_t) \propto \Pr(mx_t) f_N\left(\Delta s_t^* \mid h_t + m_{mx_t}, \upsilon_{mx_t}^2\right), \qquad t \leq T$$

where $\{m_{mx_t}, \upsilon_{mx_t}^2\}$ are the means and variances of the seven-component mixture of normal densities which are used to approximate the $\log \chi^2(1)$ distribution (see Kim, Shephard and Chib, 1998).

6. Go to step 2 and iterate $5,000$ times beyond a burn-in of $1,000$ iterations.

A.4 The Particle Filter

The particle filter of Pitt and Shephard (1999) generates a sample from the density $h_t \mid F_t, \theta$. This is a non-trivial task performed by an Auxiliary Sampling-Importance Resampling algorithm. The SV application of the algorithm is detailed in Chib, Nardari and Shephard (2002) and sketched below:

1. Given a sample $\{h^1_{t-1}, ..., h^M_{t-1}\}$ from $(h_{t-1} \mid F_{t-1}, \theta)$ calculate: $\widehat{h}^{*j}_t = \mu + \phi\left(h^j_{t-1} - \mu\right)$, $\omega_j = N\left(\Delta s_t \mid \alpha + \beta x_{t-1}, \exp(\widehat{h}^{*j}_t)\right)$, for $j = 1, ...M$. Sample $R = 10,000$ times the integers $1, 2, ..., M = 2,000$ with probability proportional to $\{\omega_j\}$. Let the sampled indices be $k_1, ..., k_R$ and associate these with $h^{*k_1}_t, ..., h^{*k_R}_t$.

2. For each value of k_j from Step 1, simulate the values $\{h^1_t, ..., h^R_t\}$ from the volatility process as: $h^{*j}_t = \mu + \phi\left(h^{k_j}_{t-1} - \mu\right) + \sigma\eta^j_t$, $j = 1, ..., R$, where $\eta^j_t \sim N(0, 1)$.

3. Resample the values $\{h^{*1}_t, ..., h^{*R}_t\}$ M times with replacement using probabilities proportional to: $\dfrac{N\left(\Delta s_t \mid \alpha + \beta_1 x_{t-1}, \exp\left(h^{*j}_t\right)\right)}{N\left(\Delta s_t \mid \alpha + \beta_1 x_{t-1}, \exp\left(h^{*k_j}_t\right)\right)}$, for $j = 1, ..., R$, to produce the desired filtered sample $\{h^1_t, ..., h^M_t\}$ from $(h_t \mid F_t, \theta)$.

References

Abhyankar, A., L. Sarno, and G. Valente (2005). "Exchange Rates and Fundamentals: Evidence on the Economic Value of Predictability," *Journal of International Economics* **66**, 325-348.

Andersen, T.G. (1996). "Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility," *Journal of Finance* **51**, 169-204.

Avramov, D. (2002). "Stock Return Predictability and Model Uncertainty," *Journal of Financial Economics* **64**, 423-458.

Backus, D.K., A.W. Gregory, and C.I. Telmer (1993). "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance* **48**, 1887-1908.

Backus, D.K., S. Foresi, and C.I. Telmer (2001). "Affine Term Structure Models and the Forward Premium Anomaly," *Journal of Finance* **56**, 279-304.

Baks, K., A. Metrick, and J.A. Wachter (2001). "Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation," *Journal of Finance* **56**, 45-85.

Bansal, R. (1997). "An Exploration of the Forward Premium Puzzle in Currency Markets," *Review of Financial Studies* **10**, 369-403.

Bansal, R., and M. Dahlquist (2000). "The Forward Premium Puzzle: Different Tales from Developed and Emerging Economies," *Journal of International Economics* **51**, 115-144.

Bansal, R., A.R. Gallant, R. Hussey, and G. Tauchen (1995). "Nonparametric Estimation of Structural Models for High Frequency Currency Market Data," *Journal of Econometrics* **66**, 251-287.

Barberis, N. (2000). "Investing for the Long Run when Returns are Predictable," *Journal of Finance* **55**, 225-264.

Bates, J.M., and C.W.J. Granger (1969). "The Combination of Forecasts," *Operations Research Quarterly* **20**, 451-468.

Bauer, G.H. (2001). "Conditional Currency Hedging and Asset Market Shocks," working paper, University of Rochester.

Bekaert, G. (1996). "The Time-Variation of Risk and Return in Foreign Exchange Markets: A General Equilibrium Perspective," *Review of Financial Studies* **9**, 427-470.

Bekaert, G., and R.J. Hodrick (1993). "On Biases in the Measurement of Foreign Exchange Risk Premiums," *Journal of International Money and Finance* **12**, 115-138.

Bekaert, G., and R.J. Hodrick (2001). "Expectations Hypotheses Tests," *Journal of Finance* **56**, 115-138.

Bekaert, G., R.J. Hodrick, and D. Marshall (1997). "The Implications of First-Order Risk Aversion for Asset Market Risk Premiums," *Journal of Monetary Economics* **40**, 3-39.

Berkowitz, J., and L. Giorgianni (2001). "Long-Horizon Exchange Rate Predictability?", *Review of Economics and Statistics* **83**, 81-91.

Bilson, J.F.O. (1981). "The 'Speculative Efficiency' Hypothesis," *Journal of Business* **54**, 435-451.

Bollerslev, T. (1986). "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics* **31**, 307-327.

Bollerslev, T. (1987). "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *Review of Economics and Statistics* **69**, 542-547.

Bollerslev, T. (1990). "Modeling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model," *Review of Economics and Statistics* **72**, 498-505.

Boudoukh, J., M. Richardson, and R.F. Whitelaw (2006). "The Information in Long-Maturity Forward Rates: Implications for Exchange Rates and the Forward Premium Anomaly," working paper, New York University.

Carter, C.K., and R. Kohn (1994). "On Gibbs Sampling for State Space Models," *Biometrika* **81**, 541-553.

Chib, S. (1995). "Marginal Likelihood from the Gibbs Output," *Journal of the American Statistical Association* **90**, 1313-1321.

Chib, S., and E. Greenberg (1995). "Understanding the Metropolis-Hastings Algorithm," *American Statistician* **49**, 327-335.

Chib, S., and I. Jeliazkov (2001). "Marginal Likelihood from the Metropolis-Hastings Output," *Journal of the American Statistical Association* **96**, 270-281.

Chib, S., F. Nardari, and N. Shephard (2002). "Markov Chain Monte Carlo Methods for Stochastic Volatility Models," *Journal of Econometrics* **108**, 281-316.

Chib, S., F. Nardari, and N. Shephard (2006). "Analysis of High Dimensional Multivariate Stochastic Volatility Models," *Journal of Econometrics* **134**, 341-371.

Clarida, R.H., L. Sarno, M.P. Taylor, and G. Valente (2003). "The Out-of-Sample Success of Term Structure Models as Exchange Rate Predictors: A Step Beyond," *Journal of International Economics* **60**, 61-83.

Clarida, R.H., L. Sarno, M.P. Taylor, and G. Valente (2006). "The Role of Asymmetries and Regime Shifts in the Term Structure of Interest Rates," *Journal of Business* **79**, 1193-1225.

Clarida, R.H., and M.P. Taylor (1997). "The Term Structure of Forward Exchange Premiums and the Forecastability of Spot Exchange Rates: Correcting the Errors," *Review of Economics and Statistics* **79**, 353-361.

Cremers, M. (2002). "Stock Return Predictability: A Bayesian Model Selection Perspective," *Review of Financial Studies* **15**, 1223-1249.

de Jong, P., and N. Shephard (1995). "The Simulation Smoother for Time Series Models," *Biometrika* **82**, 339-350.

Della Corte, P., L. Sarno, and D.L. Thornton (2007). "The Expectations Hypothesis of the Term Structure of Very Short-Term Rates: Statistical Tests and Economic Value," *Journal of Financial Economics*, forthcoming.

Diebold, F.X. (1998). "The Past, Present and Future of Macroeconomic Forecasting," *Journal of Economic Perspectives* **12**, 175-192.

Diebold, F.X. (2004). *Elements of Forecasting*, Third Edition, South-Western College Publishing.

Diebold, F.X., and P. Pauly (1990). "The Use of Prior Information in Forecast Combination," *International Journal of Forecasting* **6**, 503-508.

Engel, C. (1996). "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence," *Journal of Empirical Finance* **3**, 123-192.

Engel, C., and K.D. West (2005). "Exchange Rates and Fundamentals," *Journal of Political Economy* **113**, 485-517.

Engle, R. (2002). "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics* **20**, 339-350.

Fama, E.F. (1984) "Forward and Spot Exchange Rates," *Journal of Monetary Economics* **14**, 319-338.

Faust, J., J.H. Rogers, and J. Wright (2003). "Exchange Rate Forecasting: The Errors We've Really Made," *Journal of International Economics* **60**, 35-59.

Fernandez-Villaverde, J., and J.F. Rubio-Ramirez (2004). "Comparing Dynamic Equilibrium Models to Data: A Bayesian Approach," *Journal of Econometrics* **123**, 153-187.

Fleming, J., and C. Kirby (2003). "A Closer Look at the Relation between GARCH and Stochastic Autoregressive Volatility," *Journal of Financial Econometrics* **1**, 365-419.

Fleming, J., C. Kirby, and B. Ostdiek (2001). "The Economic Value of Volatility Timing," *Journal of Finance* **56**, 329-352.

Flood, R.B., and A.K. Rose (2002). "Uncovered Interest Parity in Crisis," *International Monetary Fund Staff Papers* **49**, 252-266.

Froot, K.A., and R.H. Thaler (1990). "Foreign Exchange," *Journal of Economic Perspectives* **4**, 179-192.

Galati, G., and M. Melvin (2004). "Why Has FX Trading Surged? Explaining the 2004 Triennial Survey," *BIS Quarterly Review*, December, 67-74.

Geweke, J. (1989). "Bayesian Inference in Econometric Models Using Monte Carlo Integration," *Econometrica* **57**, 1317-1339.

Gomez, V., and A. Maravall (2000). "Seasonal Adjustment and Signal Extraction in Economic Time Series," in Pena D., G.C. Tiao, and T.R. Tsay (eds.), *A Course in Time Series Analysis*, New York: Wiley.

Groen, J.J.J. (2000). "The Monetary Exchange Rate Model as a Long-Run Phenomenon," *Journal of International Economics* **52**, 299-319.

Hai, W., N.C. Mark, and Y. Wu (1997). "Understanding Spot and Forward Exchange Rate Regressions," *Journal of Applied Econometrics* **12**, 715-34.

Han, Y. (2006). "Asset Allocation with a High Dimensional Latent Factor Stochastic Volatility Model," *Review of Financial Studies* **19**, 237-271.

Harvey A.C., E. Ruiz, and N. Shephard (1994). "Multivariate Stochastic Variance Models," *Review of Economics Studies* **61**, 247-264.

Hlawitschka, W. (1994). "The Empirical Nature of Taylor-Series Approximations to Expected Utility," *American Economic Review* **84**, 713-719.

Hodrick, R.J. (1987). *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*, London: Harwood.

Hoeting, J.A., D. Madigan, A.E. Raftery, and C.T. Volinsky (1999). "Bayesian Model Averaging: A Tutorial," *Statistical Science* **14**, 382-417.

Kandell, S., and R.F. Stambaugh (1996). "On the Predictability of Stock Returns: An Asset-Allocation Perspective," *Journal of Finance* **51**, 385-424.

Karolyi, G.A., and R.M. Stulz (2003). "Are Financial Assets Priced Locally or Globally?", in: G. Constantinides, M. Harris, and R.M. Stulz (eds.), *Handbook of Economics of Finance*, Amsterdam: Elsevier.

Kim, S., N. Shephard, and S. Chib (1998). "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies* **65**, 361-393.

Koop, G. (2003). *Bayesian Econometrics*, New York: Wiley.

Leamer, E.E. (1978). *Specification Searches*, New York: Wiley.

Levich, R.M. (1985). "Empirical Studies of Exchange Rates: Price Behavior, Rate Determination, and Market Efficiency," in Jones, R.W., and P.B. Kenen (eds.), *Handbook of International Money and Finance*, Amsterdam: Elsevier.

Lewis, K.K. (1995). "Puzzles in International Financial Markets," in Grossman, G.M., and K. Rogoff (eds.), *Handbook of International Economics*, Amsterdam: Elsevier.

Lo, A.W. (2002). "The Statistics of Sharpe Ratios," *Financial Analysts Journal* **58**, 36-52.

Lucas, R.E. (1982). "Interest Rates and Currency Prices in a Two-Country World," *Journal of Monetary Economics* **10**, 335-359.

Lustig, H., and A. Verdelhan (2007). "The Cross-Section of Currency Risk Premia and US Consumption Growth Risk," *American Economic Review* **97**, 89-117.

Lyons, R.K. (2001). *The Microstructure Approach to Exchange Rates*, Cambridge: MIT Press.

Mark, N.C. (1995). "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review* **85**, 201-218.

Mark, N.C., and D. Sul. (2001). "Nominal Exchange Rates and Monetary Fundamentals: Evidence from a Small Post-Bretton Woods Panel," *Journal of International Economics* **53**, 29-52.

Marquering, W., and M. Verbeek (2004). "The Economic Value of Predicting Stock Index Returns and Volatility," *Journal of Financial and Quantitative Analysis* **39**, 407-429.

Meese, R.A., and K. Rogoff (1983). "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?", *Journal of International Economics* **14**, 3-24.

Obstfeld, M., and K. Rogoff (1995). "Exchange Rate Dynamics Redux," *Journal of Political Economy* **103**, 624-660.

Pitt, M.K., and N. Shephard (1999). "Filtering via Simulation Based on Auxiliary Particle Filters," *Journal of the American Statistical Association* **94**, 590-599.

Rapach, D.E., and M.E. Wohar (2002). "Testing the Monetary Model of Exchange Rate Determination: New Evidence from a Century of Data," *Journal of International Economics* **58**, 359-385.

Sarno, L., G. Valente, and H. Leon (2006). "Nonlinearity in Deviations from Uncovered Interest Parity: An Explanation of the Forward Bias Puzzle," *Review of Finance* **10**, 443-482.

Sarno, L., G. Valente, and M.E. Wohar (2004). "Monetary Fundamentals and Exchange Rate Dynamics Under Different Nominal Regimes," *Economic Inquiry* **42**, 179-193.

Shanken, J., and A. Tamayo (2001). "Risk, Mispricing, and Asset Allocation: Conditioning on Dividend Yield," NBER Working Paper 8666.

Sims, C.A. (1988). "Bayesian Skepticism on Unit Root Econometrics," *Journal of Economic Dynamics and Control* **12**, 463-474.

Tauchen, G.E., and M. Pitts (1983). "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica* **51**, 485-505.

Timmermann, A. (2006). "Forecast Combinations," in Elliott, G., C.W.J. Granger, and A. Timmermann (eds.), *Handbook of Economic Forecasting*, Amsterdam: Elsevier.

Vrontos, I.D., P. Dellaportas, and D.N. Politis (2000). "Full Bayesian Inference for GARCH and EGARCH Models," *Journal of Business and Economic Statistics* **18**, 187-198.

West, K.D., H.J. Edison, and D. Cho (1993). "A Utility-Based Comparison of Some Models of Exchange Rate Volatility," *Journal of International Economics* **35**, 23-45.

Table 1

Table 1
Descriptive Statistics for Monthly FX Returns and Fundamentals

	UK (USD/GBP)					Germany (USD/DEM − EURO)					Japan (USD/JPY)				
Panel A: Percent Returns															
	Δs_t	MF_1	MF_2	MF_3	FP	Δs_t	MF_1	MF_2	MF_3	FP	Δs_t	MF_1	MF_2	MF_3	FP
Mean	−0.012	0.000	0.000	0.000	−0.159	0.165	0.000	0.000	0.000	0.131	0.309	0.000	0.000	0.000	0.267
Std Dev	2.95	21.52	13.50	11.57	0.266	3.05	33.43	17.56	16.82	0.300	3.32	42.00	29.61	16.88	0.300
Min	−10.53	−50.02	−37.41	−24.56	−0.736	−9.71	−56.90	−34.87	−29.02	−0.779	−10.08	−72.24	−45.68	−43.25	−2.12
Max	11.69	52.89	41.86	44.00	1.35	9.21	75.85	45.75	46.78	1.59	11.89	72.29	70.28	30.85	1.31
Skewness	−0.101	−0.137	−0.372	0.474	1.49	0.156	0.164	0.240	0.473	−0.162	0.619	0.234	0.575	−0.247	−1.05
Kurtosis	4.04	2.41	3.60	4.24	9.19	3.12	2.05	2.17	2.39	4.46	4.06	1.63	2.20	2.00	13.94
$Corr\,(\Delta s_t, \Delta s_{t-1})$	0.106	0.988	0.976	0.965	0.786	0.115	0.994	0.983	0.983	0.641	0.147	0.996	0.987	0.980	0.558
$Corr\,(\Delta s_t, \Delta s_{t-3})$	0.013	0.960	0.920	0.888	0.480	0.040	0.982	0.945	0.943	0.600	0.108	0.989	0.969	0.933	0.510
$Corr\,(\Delta s_t, \Delta s_{t-6})$	0.037	0.911	0.838	0.780	0.208	0.025	0.959	0.886	0.880	0.593	−0.057	0.973	0.938	0.856	0.403
$Corr\,(\Delta s_t, \Delta s_{t-12})$	0.025	0.793	0.674	0.569	0.334	0.030	0.897	0.734	0.719	0.538	0.078	0.935	0.882	0.704	0.267
Panel B: Absolute Percent Returns															
Mean	2.25	17.92	10.46	8.57	0.237	2.43	28.14	15.07	14.18	0.271	2.48	37.17	25.65	14.60	0.315
Std Dev	1.90	11.87	8.51	7.76	0.200	1.84	17.99	8.98	9.01	0.184	2.22	19.46	14.71	8.43	0.248
Min	0.000	0.023	0.041	0.050	0.000	0.006	0.320	0.007	0.039	0.000	0.000	0.028	0.477	0.009	0.000
Max	11.69	52.89	41.86	44.00	1.35	9.71	75.85	45.75	46.78	1.59	11.89	72.29	70.28	43.25	2.12
Skewness	1.48	0.687	1.37	1.46	1.55	1.12	0.238	0.547	0.711	1.64	1.52	−0.281	0.729	0.312	1.95
Kurtosis	6.04	2.89	4.64	5.67	6.93	4.38	2.21	3.11	3.40	10.40	5.76	2.11	3.46	2.82	11.33
$Corr\,(|\Delta s_t|, |\Delta s_{t-1}|)$	0.065	0.960	0.944	0.933	0.878	0.135	0.981	0.939	0.944	0.647	0.129	0.983	0.950	0.927	0.746
$Corr\,(|\Delta s_t|, |\Delta s_{t-3}|)$	0.149	0.876	0.832	0.797	0.731	0.085	0.943	0.807	0.818	0.466	0.051	0.950	0.884	0.761	0.590
$Corr\,(|\Delta s_t|, |\Delta s_{t-6}|)$	0.093	0.724	0.690	0.606	0.577	0.051	0.876	0.633	0.641	0.410	0.010	0.882	0.786	0.541	0.500
$Corr\,(|\Delta s_t|, |\Delta s_{t-12}|)$	0.030	0.431	0.414	0.248	0.385	−0.046	0.685	0.305	0.271	0.212	−0.035	0.703	0.659	0.280	0.330

The table summarizes the descriptive statistics for the spot exchange rate percent returns (Δs_t), the three demeaned percent monetary fundamentals specifications (MF_1, MF_2, MF_3), and the percent forward premium (FP). The data sample ranges from January 1976 through December 2004 for a sample size of 348 monthly observations. The exchange rates are defined as US dollars per unit of foreign currency. For a detailed definition of the three monetary fundamentals specifications see Section 3.1.

Table 2

Table 2
Posterior Means for the UK Pound Sterling (USD/GBP)

Panel A: Bayesian Linear Regression					
Parameter	RW	MF_1	MF_2	MF_3	FP
α	-0.012	-0.012	-0.012	-0.012	-0.110
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0005)
β	$-$	0.0079	0.0254^{**}	0.0254^{*}	-0.629
		$(2.4e{-}05)$	$(3.7e{-}05)$	$(4.4e{-}05)$	(0.0016)
v^2	8.72^{***}	10.12^{***}	8.63^{***}	10.06^{***}	8.69^{***}
	(0.0021)	(0.0021)	(0.0021)	(0.0021)	(0.0021)
Panel B: Bayesian GARCH(1,1)					
α	0.018	0.027	0.005	0.017	-0.101
	(0.0022)	(0.0021)	(0.0021)	(0.0020)	(0.0022)
β	$-$	0.0042	0.0215	0.0193	-0.816
		(0.0001)	(0.0002)	(0.0002)	(0.0078)
ω	0.331^{***}	0.346^{***}	0.324^{***}	0.329^{***}	0.387^{***}
	(0.0015)	(0.0026)	(0.0017)	(0.0017)	(0.0064)
γ_1	0.905^{***}	0.902^{***}	0.903^{***}	0.902^{***}	0.897^{***}
	(0.0003)	(0.0004)	(0.0003)	(0.0003)	(0.0009)
γ_2	0.055^{***}	0.056^{***}	0.0572^{***}	0.058^{***}	0.056^{***}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0004)
Panel C: Bayesian Stochastic Volatility					
α	0.048	0.046	0.022	0.022	-0.045
	(0.0026)	(0.0028)	(0.0032)	(0.0027)	(0.0034)
β	$-$	0.0028	0.0211	0.0226	-0.653
		(0.0001)	(0.0002)	(0.0004)	(0.0107)
μ	2.01^{***}	2.02^{***}	2.01^{***}	2.01^{***}	2.00^{***}
	(0.0037)	(0.0033)	(0.0035)	(0.0036)	(0.0033)
ϕ	0.882^{***}	0.878^{***}	0.885^{***}	0.884^{***}	0.871^{***}
	(0.0012)	(0.0014)	(0.0014)	(0.0013)	(0.0015)
σ^2	0.093^{***}	0.092^{***}	0.086^{***}	0.090^{***}	0.097^{***}
	(0.0010)	(0.0010)	(0.0009)	(0.0010)	(0.0011)

The table presents the Bayesian MCMC estimates of the posterior means of the Linear Regression, GARCH(1,1) and SV model parameters for the USD/GBP monthly percent returns. The MCMC chain run for 5,000 iterations after an initial burn-in of 1,000 iterations. The numbers in parenthesis indicate the Numerical Standard Error (NSE). The superscripts *, ** and *** indicate that the 90%, 95% and 99% highest posterior density (HPD) regions, respectively, do not contain zero. The HPD region for each MCMC parameter estimate is the shortest interval that contains 95% of the posterior distribution.

<div align="center">

Table 3

Posterior Means for the Deutsche Mark/Euro (*USD/DEM-EURO*)

</div>

Panel A: Bayesian Linear Regression					
Parameter	RW	MF_1	MF_2	MF_3	FP
α	0.160 (0.0005)	0.160 (0.0005)	0.160 (0.0005)	0.160 (0.0005)	0.206 (0.0005)
β	$-$	0.0077 (1.6e−0.5)	0.0104 (3.0e−05)	0.0148 (3.1e−05)	-0.355 (0.0015)
v^2	9.30*** (0.0023)	9.26*** (0.0022)	9.30*** (0.0022)	9.27*** (0.0022)	10.81*** (0.0022)
Panel B: Bayesian GARCH(1,1)					
α	0.153 (0.0024)	0.164 (0.0022)	0.154 (0.0023)	0.159 (0.0023)	0.216 (0.0023)
β	$-$	0.0068 (7.6e−05)	0.0097 (0.0001)	0.0134 (0.0001)	-0.463 (0.0072)
ω	0.405*** (0.0019)	0.404*** (0.0026)	0.400*** (0.0028)	0.399*** (0.0027)	0.409*** (0.0023)
γ_1	0.930*** (0.0003)	0.929*** (0.0003)	0.928*** (0.0003)	0.928*** (0.0003)	0.929*** (0.0003)
γ_2	0.028*** (0.0002)	0.029*** (0.0002)	0.030*** (0.0002)	0.030*** (0.002)	0.028*** (0.0002)
Panel C: Bayesian Stochastic Volatility					
α	0.163 (0.0031)	0.176 (0.0010)	0.165 (0.0028)	0.173 (0.0028)	0.219 (0.0030)
β	$-$	0.0074 (9.4e−05)	0.0091 (0.0002)	0.0136 (0.0002)	-0.440 (0.0084)
μ	2.17*** (0.0030)	2.16*** (0.0028)	2.17*** (0.0028)	2.16*** (0.0030)	2.17*** (0.0029)
ϕ	0.746*** (0.0032)	0.751*** (0.0036)	0.757*** (0.0033)	0.755*** (0.0036)	0.751*** (0.0033)
σ^2	0.068*** (0.0005)	0.069*** (0.0006)	0.068*** (0.0005)	0.067*** (0.0005)	0.069*** (0.0005)

The table presents the Bayesian MCMC estimates of the posterior means of the Linear Regression, GARCH(1,1) and SV model parameters for the *USD/DEM-EURO* monthly percent returns. The MCMC chain run for 5,000 iterations after an initial burn-in of 1,000 iterations. The numbers in parenthesis indicate the Numerical Standard Error (NSE). The superscripts *, ** and *** indicate that the 90%, 95% and 99% highest posterior density (HPD) regions, respectively, do not contain zero. The HPD region for each MCMC parameter estimate is the shortest interval that contains 95% of the posterior distribution.

Table 4
Posterior Means for the Japanese Yen (USD/JPY)

Panel A: Bayesian Linear Regression					
Parameter	RW	MF_1	MF_2	MF_3	FP
α	0.299*	0.299*	0.299*	0.299*	0.615***
	(0.0006)	(0.0005)	(0.0005)	(0.0005)	(0.0007)
β	–	0.0070	0.0075	0.0189*	-1.224**
		(1.3e−0.5)	(1.9e−05)	(3.4e−05)	(0.0016)
v^2	11.0***	10.96***	10.99***	10.94***	10.79***
	(0.0027)	(0.0026)	(0.0026)	(0.0026)	(0.0026)
Panel B: Bayesian GARCH(1,1)					
α	0.343*	0.344*	0.342*	0.333*	0.623***
	(0.0027)	(0.0023)	(0.0024)	(0.0024)	(0.0027)
β	–	0.0065	0.0077	0.0164	-1.170**
		(6.4e−05)	(8.6e−05)	(0.0002)	(0.0069)
ω	0.595***	0.593***	0.599***	0.594***	0.676***
	(0.0028)	(0.0034)	(0.0043)	(0.0042)	(0.0076)
γ_1	0.911***	0.911***	0.909***	0.911***	0.898***
	(0.0004)	(0.0004)	(0.0005)	(0.0005)	(0.0011)
γ_2	0.037***	0.036***	0.038***	0.036***	0.041***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0007)
Panel C: Bayesian Stochastic Volatility					
α	0.166	0.166	0.161	0.138	0.532***
	(0.0039)	(0.0039)	(0.0041)	(0.0040)	(0.0038)
β	–	0.0055	0.0059	0.0155	-1.763***
		(8.1e−05)	(0.0001)	(0.0002)	(0.0109)
μ	2.16***	2.16***	2.16***	2.15***	2.06***
	(0.0040)	(0.0042)	(0.0037)	(0.0036)	(0.0040)
ϕ	0.814***	0.818***	0.818***	0.816***	0.801***
	(0.0017)	(0.0020)	(0.0018)	(0.0019)	(0.0019)
σ^2	0.149***	0.147***	0.146***	0.150***	0.230***
	(0.0025)	(0.0021)	(0.0022)	(0.0020)	(0.0043)

The table presents the Bayesian MCMC estimates of the posterior means of the Linear Regression, GARCH(1,1) and SV model parameters for the USD/JPY monthly percent returns. The MCMC chain run for 5,000 iterations after an initial burn-in of 1,000 iterations. The numbers in parenthesis indicate the Numerical Standard Error (NSE). The superscripts *, ** and *** indicate that the 90%, 95% and 99% highest posterior density (HPD) regions, respectively, do not contain zero. The HPD region for each MCMC parameter estimate is the shortest interval that contains 95% of the posterior distribution.

Table 5
The Likelihood of the Models

Panel A: In-Sample log-Likelihood

USD/GBP

Model	RW	MF_1	MF_2	MF_3	FP
LR	-867.47	-866.90	-865.12	-865.76	-866.50
$GARCH(1,1)$	-860.94	-860.51	-858.00	-858.67	-860.28
SV	-801.80	-804.07	-802.25	-801.15	-801.33

USD/DEM-EURO

LR	-878.71	-877.47	-878.09	-877.55	-878.49
$GARCH(1,1)$	-878.26	-876.96	-877.45	-876.91	-877.97
SV	-847.84	-846.23	-846.52	-846.72	-846.64

USD/JPY

LR	-907.95	-906.60	-907.17	-906.34	-904.08
$GARCH(1,1)$	-906.97	-905.60	-905.93	-905.37	-903.05
SV	-828.49	-827.76	-828.66	-827.61	-791.47

Panel B: Out-of-Sample log-Likelihood

USD/GBP

Model	RW	MF_1	MF_2	MF_3	FP
LR	-439.52	-440.173	-438.97	-439.25	-439.91
$GARCH(1,1)$	-427.58	-427.69	-426.27	-426.31	-428.25
SV	-412.08	-412.91	-412.24	-412.63	-412.18

USD/DEM-EURO

LR	-451.91	-452.37	-452.50	-453.06	-452.76
$GARCH(1,1)$	-448.47	-448.68	-449.00	-449.36	-449.03
SV	-427.120	-433.68	-433.91	-434.69	-434.58

USD/JPY

LR	-465.04	-465.58	-465.47	-466.99	-464.84
$GARCH(1,1)$	-458.33	-458.68	-458.68	-460.14	-458.05
SV	-433.814	-426.40	-425.50	-425.68	-413.52

The table reports the in-sample and out-of-sample log-likelihood values for the three FX rates (USD/GBP, USD/DEM-EURO and USD/JPY), five conditional mean specifications (RW, MF_1, MF_2, MF_3 and FP) and three volatility frameworks (Linear Regression, GARCH and Stochastic Volatility). The out-of-sample data runs from January 1990 through December 2004.

<div align="center">

Table 6

The Models with the Highest Posterior Probability

</div>

	Panel A: The Best In-Sample Models		
	Best Model	*Second Best Model*	*Third Best Model*
USD/GBP	FP^{SV}	RW^{SV}	MF_3^{SV}
$USD/DEM - EURO$	FP^{SV}	RW^{SV}	MF_2^{SV}
USD/JPY	FP^{SV}	RW^{SV}	MF_3^{SV}
	Panel B: The Best Out-of-Sample Models		
	Best Model	*Second Best Model*	*Third Best Model*
USD/GBP	RW^{SV}	FP^{SV}	MF_2^{SV}
$USD/DEM - EURO$	FP^{SV}	RW^{SV}	MF_2^{SV}
USD/JPY	FP^{SV}	RW^{SV}	MF_3^{SV}

The table shows the three best models according to the highest in-sample and out-of-sample posterior probability for the three FX rates (USD/GBP, USD/DEM-EURO and USD/JPY). The out-of-sample data runs from January 1990 through December 2004. Ranking the models using the highest posterior probability is equivalent to choosing the best model in terms of density forecasts and is a robust model selection criterion in the presence of misspecification and non-nested models.

Table 7
The Economic Value of the Empirical Exchange Rate Models

Panel A: In-Sample Performance for Models *vs.* RW^{LR}

	MF_1^{LR}				MF_2^{LR}				MF_3^{LR}				FP^{LR}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−26	−	−58	−	−129	−	−144	−	−127	−	−167	−	144	120	145	118
10%	−37	−	−90	−	−164	−	−190	−	−165	−	−228	−	180	120	181	117
12%	−51	−	−129	−	−200	−	−239	−	−205	−	−299	−	217	120	218	117

	MF_1^{GARCH}				MF_2^{GARCH}				MF_3^{GARCH}				FP^{GARCH}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	3	14	−13	−	−120	−	−128	−	−128	−	−151	−	132	101	131	98
10%	1	9	−25	−	−152	−	−165	−	−164	−	−202	−	165	101	163	96
12%	−3	−	−41	−	−184	−	−205	−	−201	−	−257	−	198	101	195	96

	MF_1^{SV}				MF_2^{SV}				MF_3^{SV}				FP^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	102	166	38	59	−55	−	−91	−	−79	−	−144	−	203	136	166	108
10%	118	151	15	18	−74	−	−132	−	−108	−	−213	−	248	132	190	97
12%	130	137	−21	−	−95	−	−180	−	−142	−	−296	−	291	129	206	87

	RW^{GARCH}					RW^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}		Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−19	−	−19	−		36	556	9	51
10%	−24	−	−24	−		42	471	16	150
12%	−28	−	−29	−		48	404	19	267

continued

Table 7 (*continued*)

Panel B: Out-of-Sample Performance for Models *vs.* RW^{LR}

	MF_1^{LR}				MF_2^{LR}				MF_3^{LR}				FP^{LR}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−173	–	−178	–	23	26	11	9	−154	–	−149	–	61	34	56	31
10%	−217	–	−224	–	27	23	9	2	−192	–	−184	–	76	34	68	30
12%	−261	–	−271	–	30	21	4	1	−229	–	−218	–	90	34	79	29

	MF_1^{GARCH}				MF_2^{GARCH}				MF_3^{GARCH}				FP^{GARCH}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−124	–	−120	–	−1	–	−2	–	−195	–	−185	–	55	28	63	32
10%	−154	–	−147	–	−1	–	−2	–	−242	–	−227	–	70	29	82	34
12%	−184	–	−172	–	−1	–	−3	–	−289	–	−267	–	86	29	103	36

	MF_1^{SV}				MF_2^{SV}				MF_3^{SV}				FP^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−18	–	−68	–	14	−3	−33	–	55	40	26	19	218	92	174	73
10%	−30	–	−108	–	10	−11	−64	–	65	37	19	10	266	90	197	65
12%	−45	–	−158	–	4	−19	−105	–	73	34	6	2	311	87	212	58

	RW^{GARCH}					RW^{SV}			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}		Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	−27	–	−17	–		105	342	76	246
10%	−32	–	−16	–		127	321	82	205
12%	−37	–	−13	–		147	304	82	166

The table presents the in-sample and out-of-sample performance fees (Φ) and break-even transaction costs (τ^{BE}) for selected models against the RW^{LR} benchmark for three target portfolio volatilities (8%, 10% and 12%). Each maximum return strategy builds an efficient portfolio by investing in the monthly return of four bonds from the US, UK, Germany and Japan and using the three exchange rates to convert the portfolio return in US dollars. The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to either 2 or 6 is willing to pay for switching from RW^{LR} to another model (such as FP^{SV}). The performance fee Φ is expressed in annual basis points. The transaction cost τ^{BE} is defined as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. The τ^{BE} values are expressed in monthly basis points and are reported only when Φ is positive. The in-sample period starts in January 1979 and the out-of-sample data runs from January 1990 through December 2004.

Table 8
The Economic Value of Combined Forecasts

Panel A: In-Sample Performance							

All Models (*vs. RW^{LR}*)							
	BMA				*BW*			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	207	145	156	120	192	117	156	93
10%	254	141	172	109	235	114	177	83
12%	299	138	178	100	276	111	191	74

Stochastic Volatility Models (*vs. RW^{LR}*)							
	BMA				*BW*			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	208	146	157	120	192	117	156	93
10%	255	142	173	110	235	114	178	83
12%	300	139	179	100	276	111	191	74

Panel B: Out-of-Sample Performance							

All Models (*vs. RW^{LR}*)							
	BMA				*BW*			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	250	130	206	108	268	128	222	107
10%	306	127	237	100	329	125	255	99
12%	360	124	259	91	386	122	279	90

Stochastic Volatility Models (*vs. RW^{LR}*)							
	BMA				*BW*			
σ_p^*	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}	Φ_2	τ_2^{BE}	Φ_6	τ_6^{BE}
8%	259	134	215	112	277	131	231	111
10%	317	131	249	104	340	129	267	103
12%	373	128	273	95	400	126	294	94

The table reports the in-sample and out-of-sample performance fees (Φ) and break-even transaction costs (τ^{BE}) for all maximum return strategies based on combined forecasts for three target portfolio volatilities (8%, 10% and 12%). *BMA* denotes *Bayesian Model Average* and *BW* is *Bayesian Winner*. The combined forecasts are shown for two cases: (i) the unrestricted universe of all 15 models, and (ii) the restricted universe of only the five stochastic volatility models. The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to either 2 or 6 is willing to pay for switching from the RW^{LR} benchmark to the *BMA* or *BW* strategy. τ^{BE} is defined as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. The transaction costs are only reported when Φ is positive. The performance fees are expressed in annual basis points, and the transaction costs in monthly basis points. The in-sample period starts in January 1979 and the out-of-sample data runs from January 1990 through December 2004.

Table 9
Out-of-Sample Robustness

Panel A: Subsample Analysis for Selected Models *vs.* RW^{LR}
$(\sigma_p^* = 10\%, \delta = 2)$

Subsample	FP^{SV}		BMA		BW	
	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}
$1990 - 1994$	40	12	196	50	260	56
$1995 - 1999$	539	347	519	346	523	357
$2000 - 2004$	229	83	227	93	224	96
$1995 - 2004$	381	193	363	208	364	281
$1990 - 2004$	266	90	306	127	329	125

Panel B: The Performance of $tGARCH$ **Models** *vs.* RW^{LR}
$(\delta = 2)$

σ_p^*	MF_1^{tGARCH}		MF_2^{tGARCH}		MF_3^{tGARCH}		FP^{tGARCH}		RW^{tGARCH}	
	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}	Φ_2	τ_2^{BE}
8%	-34	–	-32	–	-29	–	110	49	21	78
10%	-43	–	-40	–	-36	–	140	50	28	82
12%	-50	–	-48	–	-42	–	169	51	35	88

The table provides an analysis of out-of-sample robustness for the performance fees (Φ) and break-even transaction costs (τ^{BE}) of selected models against the RW^{LR} benchmark. Panel A conducts a subsample analysis and Panel B examines the performance of the tGARCH(1,1) model with Student-t innovations. *BMA* denotes *Bayesian Model Average* and *BW* is *Bayesian Winner*. All maximum return strategies build an efficient portfolio by investing in the monthly return of four bonds from the US, UK, Germany and Japan and using the three exchange rates to convert the portfolio return in US dollars. The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion equal to 2 is willing to pay for switching from RW^{LR} to (say) FP^{SV}. The target portfolio volatility in Panel A is set at 10%. τ^{BE} is defined as the minimum monthly proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy. The transaction costs are only reported when Φ is positive. The performance fees are expressed in annual basis points, and the transaction costs in monthly basis points. The combined forecasts are for the universe of all 15 models. The out-of-sample period runs from January 1990 through December 2004.

Table 10
Sharpe Ratios for Selected Models

Panel A: In-Sample				
σ_p^*	RW^{LR}	FP^{SV}	BMA	BW
8%	0.88	1.09	1.11	1.13
10%	0.91	1.14	1.15	1.17
12%	0.94	1.17	1.19	1.21
Panel B: Out-of-Sample				
σ_p^*	RW^{LR}	FP^{SV}	BMA	BW
8%	0.76	0.98	1.06	1.11
10%	0.76	0.98	1.06	1.12
12%	0.76	0.98	1.06	1.12

The table presents the in-sample and out-of-sample annualized Sharpe ratios for selected models. *BMA* denotes *Bayesian Model Average* and *BW* is *Bayesian Winner*. The Sharpe ratios are adjusted for the serial correlation in the monthly portfolio returns generated by the dynamic strategies (e.g. Lo, 2002). All maximum return strategies build an efficient portfolio by investing in the monthly return of four bonds from the US, UK, Germany and Japan and using the three exchange rates to convert the portfolio return in US dollars. The maximum return strategies are evaluated at three target portfolio return volatilities: 8%, 10%, and 12%. The in-sample period starts in January 1979 and the out-of-sample data runs from January 1990 through December 2004.
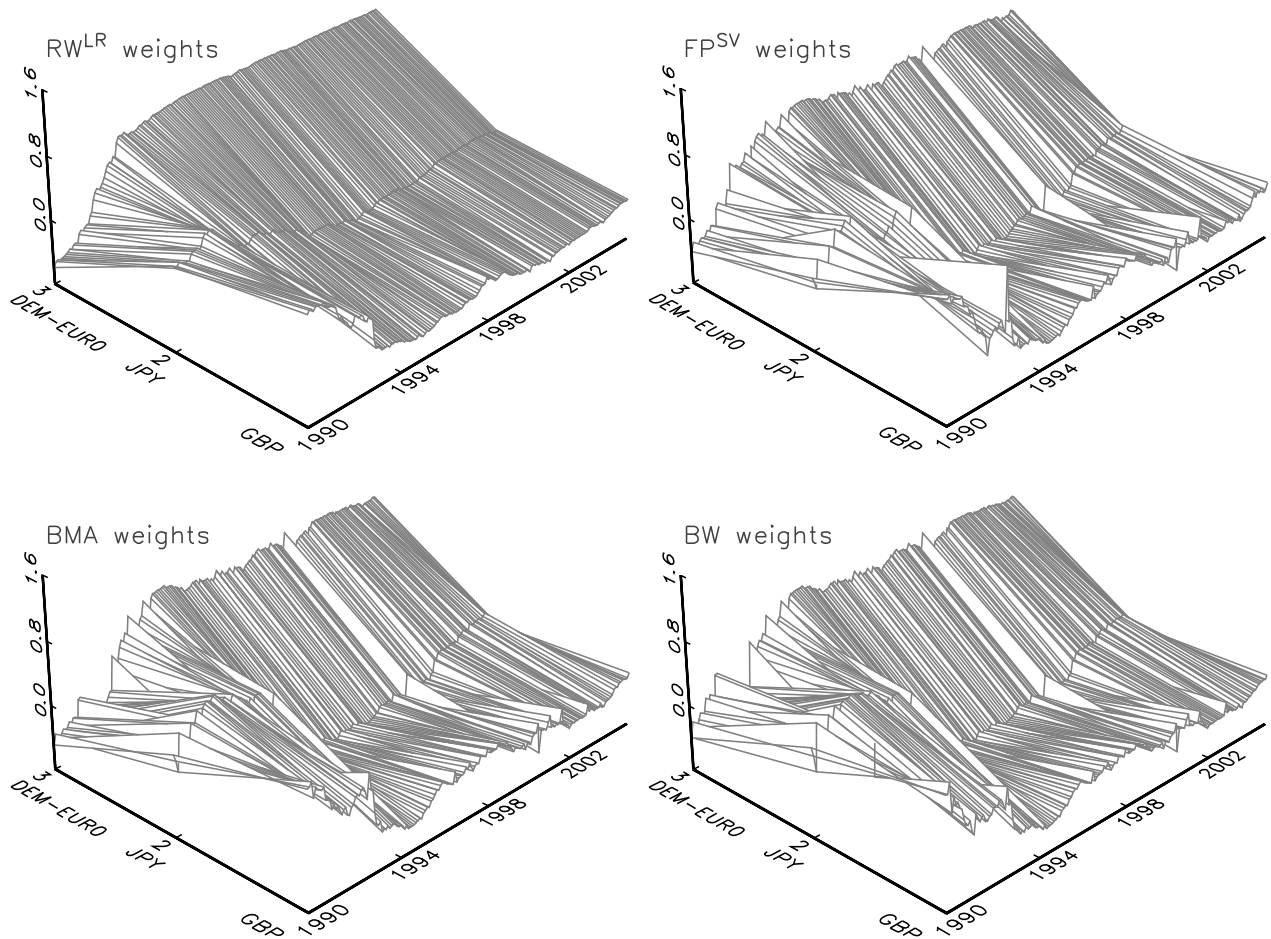
Figure 1: The Out-of-Sample Dynamic Weights for Selected Models: This is the out-of-sample time variation in the weights investing in the three risky assets (the UK, Japanese and German bonds) at a target portfolio volatility of 10% and a degree of relative risk aversion of 2. The figure presents four cases: the benchmark random walk model with constant variance (upper left), the forward premium model with stochastic volatility (upper right), the Bayesian Model Average strategy (lower left), and the Bayesian Winner strategy (lower right).