City Research Online

City St George's, University of London

Citation: Della Corte, P., Sarno, L. & Tsiakas, I. (2012). Volatility and Correlation Timing in Active Currency Management. In: James, J., Marsh, I. W. & Sarno, L. (Eds.), Handbook of Exchange Rates. (pp. 421-447). New Jersey, USA: Wiley. ISBN 9780470768839 doi: 10.1002/9781118445785.ch15

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/13162/

Link to published version: https://doi.org/10.1002/9781118445785.ch15

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to City Research Online policy.

City Research Online: http://openaccess.city.ac.uk/ publications@citystgeorges.ac.uk

Volatility and Correlation Timing in Active Currency Management*

Pasquale Della Corte
Imperial College London

Lucio Sarno
Cass Business School and CEPR

Ilias Tsiakas
University of Guelph

October 2011

Abstract

This chapter examines how dynamic volatilities and correlations in exchange rate returns affect the optimal portfolio choice of a risk-averse investor engaging in international asset allocation. We take a Bayesian approach in estimation and asset allocation that accounts for parameter and model uncertainty, and find substantial economic value in both volatility and correlation timing. This result is robust to reasonable transaction costs, parameter and model uncertainty, and alternative specifications for volatilities and correlations.

Keywords: Asset Allocation; Exchange Rates; Volatility Timing; Correlation Timing; Parameter Uncertainty; Model Uncertainty; Bayesian Model Averaging.

JEL Classification: C11; F31; F37; G11.

*__Acknowledgements:__ This paper is forthcoming as a chapter in the *Handbook of Exchange Rates*. The authors are indebted for useful conversations or constructive comments to Torben Andersen, Federico Bandi, Luc Bauwens, Valentina Corradi, Frank Diebold, Robert Engle, Gabriele Fiorentini, Rene Garcia, John Geweke, Eric Ghysels, Xing Jin, Mico Loretan, John Maheu, Anna Pavlova, Tarun Ramadorai, Neil Shephard, Georgios Skoulakis, Jun Tu, Herman van Dijk, Guofu Zhou as well as to participants at the 2008 Oxford-Man Conference in "Financial Econometrics and Vast Data" at Oxford, UK; the 2008 China International Conference in Finance at Dalian, China; the 2008 European Meeting of the Econometric Society at Milan, Italy; the 2008 Far Eastern and South Asian Meeting of the Econometric Society at Singapore; the 2008 Bayesian Workshop at the Rimini Center for Economic Analysis, Italy; the 2008 Small Open Economies in a Globalized World Conference at Waterloo, Canada; the 2007 Multivariate Volatility Models Conference at Faro, Portugal; and the CORE seminar at Universite Catholique de Louvain. *Contact details*: Pasquale Della Corte, Finance Group, Imperial College Business School, Imperial College London, 53 Prince's Gate, London SW7 2AZ, UK. Email: p.dellacorte@imperial.ac.uk; Lucio Sarno, Cass Business School, 106 Bunhill Row, London EC1Y 8TZ, UK. Email: lucio.sarno@city.ac.uk; Ilias Tsiakas, Department of Economics and Finance, University of Guelph, Guelph, Ontario N1G 2W1, Canada. Email: itsiakas@uoguelph.ca.

1 Introduction

The expected volatilities and correlations of asset returns are a critical input in the optimal portfolio choice of a risk-averse investor. Extensive empirical evidence indicates that both the volatility of asset returns and their correlations change over time.[1] However, forecasting the dynamics of volatility and correlation requires estimation of suitable multivariate models, which are notoriously complicated and difficult to handle. This has spurred a large body of empirical research exploring tractable multivariate models of time-varying volatility.[2] Among them, the Dynamic Conditional Correlation (DCC) model (Engle, 2002) has emerged as a benchmark as it provides a parsimonious and flexible framework for modeling the dynamics of asset return volatilities and correlations. Hence, it can be readily used in realistic applications of dynamic asset allocation.

This chapter addresses an essential question that lies at the core of a long line of research in empirical finance: does volatility and correlation timing matter for the optimal asset allocation of a risk-averse investor and, if so, how? We contribute to the literature on the economic value of volatility timing, which focuses primarily on the dynamics of volatilities while, however, for the most part treats the impact of dynamic correlations as an afterthought; in some cases, correlations are assumed constant (e.g., Della Corte, Sarno and Tsiakas, 2009), and in other cases they are modelled using rolling estimators (e.g., Fleming, Kirby and Ostdiek, 2001), but in no case is the effect of correlation timing separately evaluated from that of volatility timing. In this chapter, we fill this gap in the literature.

Our empirical investigation begins by estimating a large set of multivariate specifications based on the DCC model. We take a Bayesian approach in estimation and asset allocation, which allows us to evaluate volatility and correlation timing in a way that accounts for parameter and model uncertainty. Our analysis also assesses the impact of other important aspects of portfolio choice, such as transaction

[1] See, for example, Longin and Solnik (1995, 2001), Ang and Bekaert (2002), and Goetzmann, Li and Rouwenhorst (2005).

[2] For multivariate volatility models see, for example, Bollerslev, Engle and Wooldridge (1988), Diebold and Nerlove (1989), Harvey, Ruiz and Shephard (1994), Engle and Kroner (1995), Ledoit, Santa-Clara and Wolf (2003), Chib, Nardari and Shephard (2006), Philipov and Glickman (2006), Engle, Shephard and Sheppard (2008), Gourieroux, Jasiak and Sufana (2009), and Palandri (2009).

costs, asymmetry in correlations and richness of correlation dynamics.

The empirical analysis is based on 31 years of daily returns data on four major US dollar nominal exchange rates. As the largest financial market in the world, the foreign exchange (FX) market is geographically dispersed with a uniquely international dimension. Time-varying volatility in exchange rate returns is a stylized fact. Moreover, the FX market is a natural market to study dynamic correlations as investors trade currencies but all prices are quoted relative to a numeraire. For example, consider the case where the US dollar is the numeraire relative to which exchange rates are quoted. Other things being equal, a shock to the US economy will likely move the US dollar in the same direction relative to all other currencies, thus generating positive correlation in all dollar exchange rates. In general, correlations between exchange rate returns will change over time due to variation in global and country-specific fundamentals as well as other factors that are specific to the FX market, such as the intervention of policy makers aimed at influencing a particular basket of exchange rates.[3]

We assess the relative economic value of volatility and correlation timing without modeling exchange rate returns as a function of state variables.[4] This is equivalent to specifying a random walk model for the spot exchange rate, which in turn is consistent with the vast majority of the empirical FX literature since the seminal contribution of Meese and Rogoff (1983). Modeling the spot exchange rate as a random walk is also the basis of the widely used carry trade strategy that borrows in low interest rate currencies and invests in high interest rate currencies. The carry trade has historically delivered large economic gains by exploiting deviations from uncovered interest parity since on average the interest differential is not offset by a commensurate depreciation of the investment currency (e.g., Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011; Menkhoff, Sarno, Schmeling and Schrimpf, 2011).

The key distinguishing feature of our analysis is the use of economic criteria. While there is an extensive literature on statistically evaluating the performance of volatility and correlation models, there

[3]Note that transaction costs for professional investors in the FX market are very small (no more than 2 basis points), and currency hedge funds typically invest in a small number of currencies. Furthermore, in recent years investors can directly trade on FX volatility and correlation using volatility and correlation swaps (e.g., Della Corte, Sarno and Tsiakas, 2011).

[4]See Della Corte and Tsiakas (2012) in this handbook for a comprehensive statistical and economic evaluation of the predictability implied by a wide range of empirical exchange rate models for the mean.

is little work in formally assessing the economic value of volatility and correlation timing. A purely statistical analysis of volatility and correlation timing is not particularly informative to an investor as it falls short of measuring whether there are tangible economic gains from implementing dynamic conditional volatilities and correlations in active portfolio management. This motivates our asset allocation approach, which extends previous studies by West, Edison and Cho (1993) and Fleming, Kirby and Ostdiek (2001, 2003). We evaluate the dynamic allocation strategies using a constant relative risk aversion (CRRA) utility function, and measure how much a risk-averse investor is willing to pay for switching from a static portfolio strategy based on the constant covariance model to one that has dynamic conditional volatility and correlation.[5]

We assess the economic value of volatility and correlation timing in a Bayesian framework, which explicitly accounts for the fact that forecasts are not known with complete precision, and the presence of estimation error will make the resulting allocation suboptimal. The Bayesian portfolio choice literature suggests that we can account for parameter uncertainty by evaluating expected utility under the investor's predictive posterior distribution, which is determined by historical data and the prior, but does not depend on the parameter estimates (e.g., Kandel and Stambaugh, 1996; Barberis, 2000; and Kan and Zhou, 2007). We can thus examine the effect of parameter uncertainty on asset allocation by comparing, on the one hand, the standard (or "plug-in") method that replaces the true parameter values by their estimates with, on the other hand, the Bayesian approach that integrates estimation risk into the analysis.

In line with the Bayesian approach of Avramov (2002), Cremers (2002), and Della Corte, Sarno and Tsiakas (2009), we also evaluate the impact of model uncertainty by exploring whether portfolio performance improves when combining the forecasts from the large set of models we estimate.[6] Including the benchmark static covariance model, we estimate a total of 46 model specifications, which we then combine in optimally implementing Bayesian Model Averaging (BMA).

[5] For related studies, see also Engle and Sheppard (2001) and Engle and Colacito (2006).

[6] We estimate the parameters of the DCC model by designing a Markov Chain Monte Carlo (MCMC) algorithm, and thus contribute to the financial econometrics literature since the DCC model has yet to be estimated and evaluated using Bayesian methods.

Note that we carry out the empirical analysis both using the full sample of available data and by performing out-of-sample rolling estimation for selected models. Since out-of-sample estimation is very costly in terms of computation time given the current state of technology, our discussion focuses primarily on the full set of in-sample results. However, we find that the in-sample and out-of-sample results tend to be qualitatively identical.

To preview our key results, the performance of the dynamic allocation strategies suggests that there is high economic value in timing both the volatilities and the correlations of exchange rate returns. We find that an international investor facing FX risk will pay a performance fee of about 4% per year for volatility timing and a further 3% per year for correlation timing. This finding is robust to reasonable transaction costs as well as parameter uncertainty, model uncertainty, asymmetric correlations, and alternative specifications for volatilities and correlations. In particular, parameter uncertainty has little or no effect on the economic value of volatility and correlation timing. Furthermore, the simplest DCC specification examined here captures almost all gains from timing volatilities and correlations in the context of asset allocation across currencies. This is an important finding, which suggests that using a more sophisticated DCC specification will not necessarily enhance the performance of optimally designed portfolios relative to the simple DCC model.

The remainder of the chapter is organized as follows. In the next section we lay out the multivariate conditional volatility and correlation models and briefly explain the Bayesian estimation methods. Section 3 discusses the framework for assessing the economic value of volatility and correlation timing for a risk-averse investor with a CRRA portfolio allocation strategy. The effect of parameter uncertainty on asset allocation is discussed in Section 4, while model uncertainty and the construction of combined forecasts are described in Section 5. Our empirical results are reported in Section 6 and Section 7 concludes. We also include Appendix A with details on the volatility models employed in this chapter and Appendix B with further details on how we account for parameter uncertainty.

2 Dynamic Models for Volatility and Correlation

We model the dynamics of volatilities and correlations of exchange rate returns using a set of specifications based on the Dynamic Conditional Correlation (DCC) model (Engle, 2002). The DCC model offers an attractive multivariate framework for the study of correlation timing because it has the following advantages: it is tractable and parsimonious with a low dimension of parameters; it is flexible and can be generalized to account for asymmetric correlations, while ensuring that correlations are in the $[-1, 1]$ range; it provides for a positive-definite covariance matrix; and, finally, it is straightforward to implement even when the number of assets is large.

In order to assess the economic value of volatility and correlation timing, we estimate a set of multivariate models for dynamic correlations (such as the DCC model), each under a set of univariate specifications for dynamic volatility (such as the GARCH model). In what follows we describe the complete set of models we estimate.

2.1 The Set of Multivariate Models

Let $y_t = (y_{1,t}, \ldots, y_{N,t})'$ denote the $N \times 1$ vector of nominal log-exchange rate returns at time t:

$$y_t = \mu + \Sigma_t^{1/2} \varepsilon_t, \tag{1}$$

where $\mu = (\mu_1, \ldots, \mu_N)'$ is the $N \times 1$ vector of unconditional means, Σ_t is the $N \times N$ conditional covariance matrix, and $\varepsilon_t = (\varepsilon_{1,t}, \ldots, \varepsilon_{N,t})'$ is the $N \times 1$ vector of standard normal disturbances.[7]

In our analysis the vector of means μ is constant over time. This is equivalent to specifying a random walk model for each log-exchange rate. This assumption allows us to primarily focus on the effect of dynamic volatility and correlation on asset allocation. Note that the predictability in the mean implied by a wide range of empirical exchange rate models is studied separately by Della Corte and Tsiakas (2012) in this handbook. For the conditional covariance matrix Σ_t, we use the set of specifications described below.

[7] See Bauwens and Laurent (2005) for a generalization of the conditional normality assumption of multivariate models.

2.1.1 The Static Benchmark

The *Multivariate Linear Regression* (MLR) model assumes a constant covariance matrix and represents our benchmark model. This corresponds to setting $\Sigma_t = \overline{\Sigma}$. The MLR model can be viewed as the carry trade strategy widely adopted in currency markets (e.g., Burnside *et al.*, 2011). The alternative models presented below reflect strategies that augment the carry trade with volatility and correlation timing.

2.1.2 The Constant Conditional Correlation Model

The *Constant Conditional Correlation* (CCC) model (Bollerslev, 1990) assumes constant correlations but dynamic volatilities. This model decomposes the conditional covariance matrix as follows:

$$\Sigma_t = D_t \overline{R} D_t, \tag{2}$$

$$D_t = diag\left\{\sigma_{1,t}, \ldots, \sigma_{N,t}\right\}, \tag{3}$$

where D_t is the $N \times N$ diagonal matrix of dynamic conditional volatilities $\{\sigma_{j,t}\}$ for $j \leq N$, and $\overline{R}$ is the $N \times N$ matrix of unconditional correlations. The conditional volatilities can have any of the specifications discussed in Section 2.2 below. The main feature of the CCC model is that the dynamics of covariances are governed exclusively by the dynamics of volatilities since correlations are constant.

2.1.3 The Dynamic Conditional Correlation Model

The *Dynamic Conditional Correlation* (DCC) model (Engle, 2002) assumes dynamic volatilities and correlations by decomposing the conditional covariance matrix as follows:

$$\Sigma_t = D_t R_t D_t, \tag{4}$$

$$D_t = diag\left\{\sigma_{1,t}, \ldots, \sigma_{N,t}\right\}, \tag{5}$$

$$R_t = \left(diag\left\{Q_t\right\}\right)^{-1/2} Q_t \left(diag\left\{Q_t\right\}\right)^{-1/2}, \tag{6}$$

$$Q_t = (\overline{R} - \Gamma' \overline{R} \Gamma - \Delta' \overline{R} \Delta) + \Gamma' z_{t-1} z_{t-1}' \Gamma + \Delta' Q_{t-1} \Delta, \tag{7}$$

where D_t is the $N \times N$ diagonal matrix of dynamic conditional volatilities, R_t is the $N \times N$ symmetric matrix of dynamic conditional correlations, $\overline{R}$ is the $N \times N$ matrix of unconditional correlations, Q_t is an $N \times N$ symmetric positive-definite matrix, Γ and Δ are $N \times N$ parameter matrices, and $z_t = D_t^{-1} u_t \sim N(0, R_t)$, where $u_t = y_t - \mu$.

We focus on two versions of this model. The simplest version reduces $\Gamma = \gamma$ and $\Delta = \delta$, where $\{\gamma, \delta\}$ are scalars, which are the same for all assets $i \leq N$. This is the "scalar" *DCC* model, which assumes that the dynamics of all correlations are driven by the same two parameters $\{\gamma, \delta\}$. A less parsimonious variant of the scalar *DCC* model results when the matrices Γ and Δ are assumed to be diagonal: $\Gamma = diag\{\gamma_1, ..., \gamma_N\}$ and $\Delta = diag\{\delta_1, ..., \delta_N\}$. This is the "diagonal" *DCC* model, which allows for distinct dynamics in each correlation process but requires estimation of more parameters. We estimate both scalar (henceforth denoted simply as *DCC*) and diagonal (DCC_{diag}) specifications.

2.1.4 The Asymmetric Dynamic Conditional Correlation Model

The *Asymmetric Dynamic Conditional Correlation* (*ADCC*) model (Cappiello, Engle and Sheppard, 2006) further allows for asymmetric correlations by generalizing Equation (7) as follows:

$$Q_t = (\overline{R} - \Gamma'\overline{R}\Gamma - \Delta'\overline{R}\Delta - \Pi'\overline{P}\Pi) + \Gamma' z_{t-1} z_{t-1}'\Gamma + \Delta' Q_{t-1}\Delta + \Pi' p_{t-1} p_{t-1}'\Pi, \tag{8}$$

where Γ, Δ and Π are $N \times N$ parameter matrices; $p_t = I[z_t < 0] \circ z_t$, where $I[\cdot]$ is an indicator function taking the value 1 if the argument is true and 0 otherwise, and $\circ$ indicates the Hadamard product; and $\overline{P} = E[p_t p_t']$. For example, the symmetric scalar *DCC* model (the simplest *DCC* model we consider) is obtained as a special case of the *ADCC* model when $\Gamma = \gamma$, $\Delta = \delta$ and $\Pi = 0$. As with the symmetric version, we estimate both scalar (*ADCC*) and diagonal ($ADCC_{diag}$) specifications.

The *ADCC* model is motivated by numerous empirical studies showing that return correlations may be asymmetric as they tend to increase in highly volatile bear markets (e.g., Longin and Solnik, 2001; Ang and Chen, 2002). This has important implications for optimal asset allocation and, for example, casts doubt on the benefits of international diversification (e.g., Ang and Bekaert, 2002). When asset returns are more volatile, investors have a stronger incentive to diversify, but it is precisely in these

cases that correlations are high and diversification opportunities are low. In other words, asymmetric return correlations cause diversification opportunities to be least available when they are most needed.[8] In this context, the *ADCC* model allows us to determine the possible impact of asymmetric correlations on asset allocation in the FX market.

2.2 The Set of Univariate Models for Volatility Timing

We estimate the multivariate models under a variety of univariate specifications for the conditional variance, including some of the most popular models in the literature. The nine univariate volatility models we consider are: GARCH (Bollerslev, 1986); AVGARCH (Taylor, 1986); NARCH (Higgins and Bera, 1992); EGARCH (Nelson, 1991); ZARCH (Zakoian, 1994); GJR-GARCH (Glosten, Jagannathan and Runkle, 1993); APGARCH (Ding, Engle and Granger, 1993); AGARCH (Engle, 1990); and NAGARCH (Engle and Ng, 1993). The details on the full specification of these models are available in Appendix A.

2.3 Pairwise Model Comparisons

In addition to the static benchmark MLR, the set of models comprises five competing dynamic specifications $\{CCC, DCC, DCC_{diag}, ADCC \text{ and } ADCC_{diag}\}$ under each of the nine univariate volatility specifications listed above. In total, therefore, we estimate 46 model specifications. The principal objective of our analysis is to provide an economic evaluation of these models in the context of dynamic asset allocation. First, we assess the economic value of volatility timing simply by comparing the CCC model to the static MLR. We then measure the additional economic gains from correlation timing by comparing the DCC to the CCC model. We also assess whether there is economic value in imposing separate dynamics on correlations (diagonal DCC vs. scalar DCC), whether correlation asymmetries are important ($ADCC$ vs. DCC), and finally, whether the choice of a particular volatility specification generates further economic gains.

[8] Recent work in empirical asset pricing shows that asymmetric correlation risk is priced in the sense that assets which pay off well when market-wide correlations are higher than expected earn negative excess returns. The negative excess return on correlation-sensitive assets can therefore be interpreted as an insurance premium (e.g., Driessen, Maenhout and Vilkov, 2009; and Buraschi, Porchia and Trojani, 2010).

2.4 Estimation and Forecasting

We perform Bayesian estimation of all model parameters. The critical advantage of the Bayesian methodology is that it provides a unified framework for estimation, forecasting and model selection, which is particularly suitable for capturing the effect of parameter and model uncertainty. Bayesian inference generally provides the posterior distribution of the parameters conditional on the data, which holds for finite samples. The posterior distribution can in turn be used as an input in forming Bayesian asset allocation strategies for an economically meaningful ranking of the models that accounts for parameter and model uncertainty.

In this chapter, we design a Markov Chain Monte Carlo (MCMC) algorithm for Bayesian estimation of the *DCC* model. The algorithm draws insights from the Bayesian stochastic volatility algorithm of Kim, Shephard, and Chib (1998), Chib, Nardari, and Shephard (2002, 2006) and Tsiakas (2006), and from the Bayesian univariate *GARCH* algorithm of Della Corte, Sarno and Tsiakas (2009). The Bayesian MCMC algorithm constructs a Markov chain whose limiting distribution is the target posterior density of the 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. Our Bayesian estimation approach delivers a sample from the posterior distribution of the parameters, which is a key input to the Bayesian asset allocation.[9]

3 The Economic Value of Volatility and Correlation Timing

3.1 The Dynamic Strategy

We design an international asset allocation strategy that involves trading the US dollar and four other currencies: the British pound, Deutsche mark/euro, Swiss franc and Japanese yen. Consider a US

[9] A detailed description of the MCMC algorithm is available from the authors upon request. The algorithm produces estimates of the posterior means of $\theta = \{\mu, \theta_1, \theta_2\}$, where $\mu = \{\mu_i\}$ is the vector of unconditional means for assets $i \leq N$, θ_1 are the parameters of each univariate GARCH-type volatility process, and θ_2 are the correlation parameters. For example, for the diagonal $ADCC_{diag}$ model with $GARCH$ volatility: $\theta_1 = \{\omega_i, \alpha_i, \beta_i\}$ and $\theta_2 = \{\gamma_i, \delta_i, \pi_i\}$. Setting $N = 4$ requires 28 parameter estimates.

investor who builds a portfolio by allocating her wealth between five bonds: one domestic (US), and four foreign bonds (UK, Germany, Switzerland and Japan). The yield of the bonds is proxied by eurodeposit rates. At the beginning of each period, the foreign bonds yield a riskless return in local currency but a risky return in US dollars. Hence the only risk the US investor is exposed to is FX risk. Every period the investor takes two steps. First, she uses each model to forecast the one-day ahead conditional volatilities and correlations of the exchange rate returns. Second, conditional on the forecasts of each model, she dynamically rebalances her portfolio by computing the new optimal weights that maximize utility. This setup is designed to inform us whether using one particular conditional volatility and correlation specification affects the performance of an allocation strategy in an economically meaningful way.

3.2 Dynamic Asset Allocation with CRRA Utility

We set up a dynamic asset allocation framework with Constant Relative Risk Aversion (CRRA) utility for assessing the economic value of strategies that exploit predictability in volatilities and correlations. Consider the portfolio choice at time t of an investor who maximizes the expected end-of-period utility by trading in every period N risky assets and a risk-free asset. The investor problem is formally defined as follows (e.g., Brandt, 2009):

$$\begin{aligned} V_t\left(W_t, Z_t\right) &= \max_{\{x_s\}_{s=t}^{T-1}} E_t\left[\frac{W_T^{1-\lambda}}{1-\lambda}\right] \\ s.t.\ \ W_{s+1} &= W_s R_{p,s+1} \quad \forall s \geq t, \end{aligned} \tag{9}$$

where W_T is the end-of-period wealth, λ denotes the coefficient of relative risk aversion (RRA), $R_{p,s+1} = x_s'\left(R_{s+1} - R^f\right) + R^f$ is the gross portfolio return from time s to $s+1$, x_s is the vector of portfolio weights on the risky assets chosen at time s, R_{s+1} is the vector of gross returns on the risky assets from time s to $s+1$, R^f is the gross return on the risk-free asset, and Z_t is the information set available at time t captured by the set of conditional volatilities and correlations.[10] In this intertemporal portfolio problem, at date t the investor optimally chooses the portfolio weights x_t conditional on having wealth

[10] Note that we allow the risk-free rate to vary over time but we drop the time subscript for notational simplicity.

W_t and information Z_t, while taking into account that at every future date s the portfolio weights will be optimally revised conditional on the then available wealth W_s and information Z_s. The value function $V_t(W_t, Z_t)$ denotes the expectation, conditional on information Z_t, of the utility of terminal wealth W_T generated by current wealth W_t and the optimal portfolio weights $\{x_s^*\}_{s=t}^{T-1}$. Following Marquering and Verbeek (2004) and Han (2006) we set $\lambda = 6$, which produces portfolios with reasonable expected return and volatility.

This intertemporal allocation problem does not have a simple and tractable solution as would be the case in a mean-variance setting. We solve for the optimal portfolio choice using the method developed by Brandt, Goyal, Santa-Clara and Stroud (2005). This is a simulation-based method that allows for non-standard preferences, a large number of state variables and a large number of assets with arbitrary return distributions. More importantly, the method allows us to directly use the volatility and correlation forecasts from the models in computing the dynamic weights as well as extend our analysis to a Bayesian setting where expected utility is evaluated under the posterior predictive density.

3.3 Performance Measures

Following Fleming, Kirby and Ostdiek (2001), we measure the economic value of volatility and correlation timing using the principle that, at any point in time, one set of estimates of the conditional volatilities and correlations is better than a second set if investment decisions based on the first set lead to higher utility. We compute the performance fee, Φ, by equating the average utility of the MLR optimal portfolio with the average utility of the competing optimal portfolio (say the DCC portfolio) that is subject to daily expenses Φ. Since the investor is indifferent between these two strategies, we interpret Φ as the maximum performance fee she will pay to switch from the MLR to the DCC strategy. In other words, this utility-based criterion measures how much a CRRA investor is willing to pay for conditioning on dynamic volatility and correlation forecasts. To estimate the performance fee, we find the value of Φ that satisfies:

$$\sum_{t=0}^{T-1} E_t\left[U\left(R_{p,t+1}^* - \Phi\right)\right] = \sum_{t=0}^{T-1} E_t\left[U\left(R_{p,t+1}\right)\right], \tag{10}$$

where $U(\cdot)$ is the CRRA utility function, $R^*_{p,t+1}$ is the gross portfolio return constructed using the expected return, volatility and correlation forecasts from the DCC model (or another competing model), and $R_{p,t+1}$ is the gross portfolio return implied by the benchmark MLR model. We report Φ in annualized basis points.

3.4 Transaction Costs

The effect of transaction costs is an essential consideration in assessing the profitability of trading strategies. We account for this effect by calculating the break-even proportional transaction cost, τ^{be}, that renders investors indifferent between two strategies (e.g., Han, 2006). In comparing a dynamic strategy with the benchmark MLR strategy, an investor who pays transaction costs lower than τ^{be} will prefer the dynamic strategy. Since τ^{be} is a proportional cost paid every time the portfolio is rebalanced, we report τ^{be} in daily basis points.[11]

4 Parameter Uncertainty in Bayesian Asset Allocation

The asset allocation literature typically assumes that investors make optimal decisions with full knowledge of the true parameters of the model. In practice, model parameters have to be estimated, and if there is estimation error the resulting allocation will be suboptimal. This gives rise to estimation risk in applications of the plug-in method that replaces the true parameter values by their estimates. In contrast, the Bayesian approach to asset allocation integrates estimation risk into the analysis and deals with parameter uncertainty by assuming that the investor evaluates her expected utility under the predictive distribution, which is determined by historical data and the prior, but does not depend on the parameter estimates.

We consider an investor who takes into account volatility and correlation timing but is uncertain about the parameters of the model. The Bayesian portfolio choice literature argues that in the presence of parameter uncertainty, the unknown objective return distribution in the expected utility maximization

[11] In recent years, the typical transaction cost a large investor pays in the FX market is 1 pip, which is equal to 0.01 cent. For example, if the USD/GBP exchange rate is equal to 1.5000, 1 pip would raise it to 1.5001 and this would roughly correspond to 1/2 basis point proportional cost.

should be replaced with the investor's subjective posterior return distribution reflecting the information contained in the historical data and the investor's prior beliefs about the parameters. Using predictive distributions was pioneered by Zellner and Chetty (1965) and used, among others, by Kandel and Stambaugh (1996), Barberis (2000), and Kan and Zhou (2007). These studies demonstrate that parameter uncertainty is an important dimension of risk, which can substantially affect the investor's optimal allocation. For example, Barberis (2000) shows that for a long-run investor the Bayesian solution is more conservative than the plug-in approach by taking smaller positions in the risky assets. Intuitively, the Bayesian approach explicitly recognizes estimation risk as an additional source of risk, and therefore, the riskless asset becomes a more attractive investment.

In short, our plan for understanding the effect of parameter uncertainty on volatility and correlation timing is to compare, on the one hand, the allocation of an investor who uses the predictive distribution in forecasting volatilities and correlations with, on the other hand, the allocation of an investor who ignores estimation error, sampling instead from the distribution of returns conditional on fixed parameter estimates. See Appendix B for more details.

5 Model Uncertainty

Model uncertainty (or model risk) arises from the uncertainty over selecting a model specification. Consistent with our Bayesian approach, a natural criterion for resolving this uncertainty is to construct combined forecasts based on the posterior probability of each model. The posterior probability has three important advantages: (i) it is based on the marginal likelihood, and therefore accounts for parameter uncertainty,[12] (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. In computing the posterior probabilities, we set our prior belief to be that all models are equally likely.

We construct combined forecasts based on the Bayesian Model Average (BMA) strategy and the Bayesian Model Winner (BMW) strategy (e.g., Geweke and Whiteman, 2006; and Timmermann, 2006).

[12] Note that the marginal likelihood is an averaged (not a maximized) likelihood, and hence it integrates out parameter uncertainty. We compute the marginal likelihood as in Chib and Jeliazkov (2001).

In assessing the economic value of combined forecasts, we treat the BMA and BMW strategies the same way as any of the individual models. For instance, we compute the performance fee and the break-even transaction cost for the BMA relative to the MLR benchmark. We apply BMA and BMW to three universes of models: (i) VOL is the universe of all nine univariate volatility specifications under the scalar DCC; (ii) $CORR$ is the universe of the five multivariate correlation specifications (CCC, DCC, DCC_{diag}, $ADCC$, $ADCC_{diag}$) with $GARCH$ volatility; and (iii) $FULL$ is the complete universe of all 46 models, including the benchmark MLR.

5.1 The BMA Strategy

The BMA strategy accounts directly for uncertainty in model selection, and is straightforward to implement once we have the output from the MCMC simulations (e.g., Wright, 2008). The BMA volatility and correlation forecasts are simply a weighted average of the volatility and correlation forecasts across the K competing models using as weights the posterior probability of each model. Note that the BMA strategy is evaluated ex ante as the weights are set at time t and the forecasts are for time $t+1$.

5.2 The BMW Strategy

Under the BMW strategy, in each time period we select the set of one-step ahead conditional volatilities and correlations from the model that has the highest posterior probability in that period. In other words, every period the BMW strategy only uses the forecasts of the "winner" model in terms of posterior probability, and hence discards the forecasts of the other models. Clearly, there is no model averaging in the BMW strategy. Similar to the BMA, the BMW strategy is evaluated ex ante.

6 Empirical Results

6.1 Data and Descriptive Statistics

Our analysis employs daily returns data on four major US dollar nominal spot exchange rates over the period of January 1976 to December 2006 corresponding to a total of $8,069$ observations. The exchange rates are taken from *Datastream* and include the British pound (GBP), the Deutsche mark/euro (EUR), the Swiss franc (CHF) and the Japanese yen (JPY). After the introduction of the euro in January 1999,

we use the official Deutsche mark/euro conversion rate to obtain the EUR series. The exchange rate is defined as the US dollar price of a unit of foreign currency so that an increase in the exchange rate implies a depreciation of the US dollar.

Table 1 reports descriptive statistics for the daily percent exchange rate returns. For our sample period, the means are near zero ranging from -0.0004 (or -0.1% per annum) for GBP to 0.0117 (or 2.9% per annum) for JPY. The daily standard deviations revolve between 0.620 for GBP (or 9.8% per annum) to 0.736 for CHF (or 11.7% per annum). Skewness is negative for two of the four FX returns, while kurtosis ranges from 6.02 for EUR to 9.78 for GBP. Finally, the average return cross-correlations are strongly positive ranging between 0.336 and 0.819.

6.2 Bayesian Estimation

We begin by performing Bayesian estimation of the parameters of all models set out in Section 2. In addition to the MLR benchmark, the universe of models includes another 45 specifications: CCC, DCC, DCC_{diag}, $ADCC$ and $ADCC_{diag}$, each under nine alternative GARCH-type volatility specifications.

Table 2 presents the posterior mean estimates for the parameters of the asymmetric diagonal $ADCC_{diag}$ model with $GARCH$ innovations. We only report the parameter estimates of this model because it is the most general specification we consider in our analysis. The table shows that both volatilities and correlations are highly persistent for all four FX return series. This is a first indication that volatilities and correlations are predictable. Furthermore, the parameters $\{\pi_i\}$ indicating asymmetry in dynamic correlations are small for all exchange rates, with the exception of the British pound.

In our Bayesian framework, we assess statistical significance by reporting 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 also report the Numerical Standard Error (NSE) of each parameter, which provides a measure of convergence in the MCMC chain. Table 2 shows that statistical significance is much stronger for the volatility and correlation persistence parameters

than for the unconditional means. However, all parameter estimates exhibit very low NSE values, thus indicating that the estimates across the MCMC iterations have converged to the posterior means.

With the parameter estimates at hand, we generate the volatility and correlation forecasts used in the asset allocation. Figure 1 illustrates the daily in-sample correlation forecasts from the simple $DCC-GARCH$ model.

6.3 Evaluating Volatility and Correlation Timing

We assess the economic value of FX volatility and correlation timing by analyzing the performance of dynamically rebalanced portfolios constructed using the set of forecasts from the multivariate models. Our forecasting analysis is conducted both in sample and out of sample. We first discuss the in-sample results for which we have the full range of empirical findings. The out-of-sample analysis is more focused than the in-sample analysis because it is computationally very demanding to re-estimate period-by-period 46 multivariate model specifications over a 31-year sample. The out-of-sample results are reviewed at the end.

The economic evaluation focuses on the performance fee that a US investor is willing to pay for switching from the benchmark to a competing strategy. The in-sample fees are reported in Table 3, which shows the economic value of each volatility and correlation specification relative to the benchmark MLR model. The table first reports the results for CRRA utility under the plug-in method for an extensive set of models; then, for a smaller number of selected models we allow for parameter uncertainty only in volatilities and correlations, but not in the mean; finally, we present the results with full parameter uncertainty in the mean, volatilities and correlations.

We first show that there is substantial economic value associated with volatility timing. Then we show that there is high economic value in timing dynamic FX correlations over and above the economic value of volatility timing. We illustrate these results by analyzing portfolio performance: switching from the static MLR to the $CCC-GARCH$ model gives a high fee of 386 annual basis points (bps); switching from MLR to the $DCC-GARCH$ model raises the fee to 696 bps. This is also reflected in the Sharpe

ratio (SR) of the strategies: the SR rises from 1.08 for the MLR to 1.36 for $CCC-GARCH$, and then to 1.62 for $DCC-GARCH$. In short, these results demonstrate that it is worth using high-dimensional multivariate models for dynamic correlations as they generate significant economic value.[13]

Table 3 reveals another important finding that justifies the choice of economic criteria in assessing alternative specifications for dynamic correlations. A purely statistical analysis may conclude that the rich correlation structure of diagonal DCC models with asymmetric correlations leads to improved performance (e.g., Cappiello, Engle and Sheppard, 2006).[14] This is not the case in evaluating correlation timing in the context of asset allocation as the choice of dynamic volatility specification (e.g., $GARCH$ vs. $EGARCH$) or dynamic correlation specification (e.g., DCC vs. $ADCC$) has little effect on the results. We find that economic value is generated by making volatilities and correlations dynamic, irrespective of their exact specification. Therefore, not only is the investor much better off with dynamic correlations, but the simple (scalar symmetric) DCC model with $GARCH$ volatility is as good as any other model we consider. In our framework, increasing the sophistication of the econometric specification does not enhance the economic value of the simple DCC model.

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 benchmark. We evaluate the impact of transaction costs on dynamic asset allocation by computing the break-even transaction cost, τ^{be}, as the minimum proportional cost that cancels out the utility advantage (and hence positive performance fee) of a given strategy. In comparing a dynamic strategy with the static MLR strategy, an investor who pays a transaction cost lower than τ^{be} will prefer the dynamic strategy. The τ^{be} values are expressed in daily basis points. Table 3 shows that the τ^{be} values generally revolve around 10 *bps* for constant correlation models and 11 *bps* for dynamic correlation models. Given that the cost of portfolio rebalancing for large investors in the FX market is around 1 or 2 *bps*, we conclude

[13] Note that the good performance of the static MLR benchmark primarily reflects the profitability of the standard carry trade strategy and the fact that these results are in sample. As we will see later, the out-of-sample Sharpe ratios tend to be lower but confirm that dynamic volatilities and correlations significantly enhance the performance of dynamically rebalanced FX portfolios.

[14] We can show this by computing the log-likelihood values or posterior model probabilities. Since we focus on the economic value of the models, we do not report these statistical results, but they are available upon request.

that the economic value of correlation timing is robust to reasonable transaction costs.[15]

6.3.1 The Effect of Parameter Uncertainty

The plug-in approach we have discussed so far takes the parameter estimates as true and ignores estimation error. Intuitively, however, the more parameters we estimate the more uncertain we are about the validity of our volatility and correlation forecasts. We address this concern by evaluating expected utility under the Bayesian predictive density, which reflects the posterior information contained in the returns data and the investor's prior beliefs, but does not depend on the parameter estimates. Table 3 shows that parameter uncertainty in second moments has little or no effect on portfolio performance. This result has a simple intuitive explanation. Volatilities and correlations are highly persistent and tend to be estimated with high precision. As a consequence, parameter uncertainty in volatilities and correlations will not play a prominent role for a one-step ahead predictive horizon.

In contrast, parameter uncertainty turns out to be more important for the first moments, which are notoriously difficult to estimate with high precision. Our results suggest that the unconditional means in all models are hard to pin-point with great accuracy. For example, the fee for the *DCC* model falls from 696 in the plug-in case to 655 in the full parameter uncertainty case, whereas the Sharpe ratio falls from 1.62 to 1.58. As discussed in Abhyankar, Sarno and Valente (2005), parameter uncertainty makes FX investors more cautious by taking less risk and thus expecting a lower reward-to-risk ratio. Since parameter uncertainty is an additional source of risk in asset allocation, investors will optimally choose less risky portfolios, and thus attain lower fees and lower Sharpe ratios. At any rate, despite the uncertainty over a large number of parameters, the economic gains from volatility and correlation timing remain strong.

[15] We also calculate the portfolio performance of selected models using mean-variance utility since this is what is predominantly used by the literature (e.g., West, Edison and Cho, 1993; Fleming, Kirby and Ostdiek, 2001; Della Corte, Sarno and Thornton, 2008; Della Corte, Sarno and Sestieri, 2011). We find that the economic value of correlation timing with mean-variance utility is still high but with CRRA utility it tends to be a bit more pronounced as the latter captures the effect of higher order moments. These results are available upon request.

6.3.2 The Effect of Model Uncertainty

We evaluate the effect of model uncertainty on volatility and correlation timing by exploring whether portfolio performance improves when combining the forecasts from the large set of models we estimate. We focus on the Bayesian Model Averaging (BMA) and Bayesian Model Winner (BMW) strategies, which are evaluated ex ante and are applied to three universes of models: VOL is the universe of all nine $GARCH$-type univariate volatility specifications we estimate under the simple (scalar symmetric) DCC model; $CORR$ is the universe of the five multivariate correlation specifications (CCC, DCC, DCC_{diag}, $ADCC$, $ADCC_{diag}$,) with $GARCH$ volatility; and $FULL$ is the complete universe of all 46 models (including the benchmark MLR).

The economic value of combined forecasts is reported in Table 4, which assesses the impact of model uncertainty under the Bayesian predictive density. Hence we evaluate model uncertainty while at the same time we account for parameter uncertainty. The results in Table 4 indicate that there is high economic value in BMA forecast combinations and even higher value in the BMW strategy. For instance, compared to the $DCC-GARCH$ model, which under full parameter uncertainty delivers a Sharpe ratio of 1.58, the $BMA-FULL$ combined strategy increases the Sharpe ratio to 1.73 and the $BMW-FULL$ raises it further to 1.81. In conclusion, accounting for model uncertainty by forming combined forecasts delivers additional economic value and makes the case for volatility and correlation timing more robust.

6.3.3 Out-of-Sample Portfolio Performance

The results reported so far are in sample for the period of January 1976 to December 2006. In this section we discuss out-of-sample portfolio performance. The out-of-sample period ranges from January 1986 to December 2006 and uses a rolling window of 10 years. Since this exercise is very computationally intensive, we generate daily out-of-sample forecasts using month-by-month parameter estimates for the means, volatilities and correlations. This means, for instance, that all daily forecasts generated for January 1986 are based on parameters estimated using information up to the end of December

1985. Moreover, we focus our out-of-sample discussion on the family of DCC models with $GARCH$ innovations since $GARCH$ volatility performs as well as any other volatility specification.

The results are reported in Table 5 and show that the out-of-sample findings are qualitatively similar to the in-sample findings. The out-of-sample Sharpe ratios confirm that there is still high incremental economic value in both volatility and correlation timing: the Sharpe ratio is 0.57 for the MLR, 0.86 for the CCC and 0.96 for the DCC. The out-of-sample Sharpe ratios are plotted in Figure 2. The performance fees are of the same order of magnitude as in sample: 457 *bps* for the CCC and 800 *bps* for the DCC. Again, diagonal or asymmetric DCC specifications do not add further economic value to the simple scalar symmetric DCC. Finally, the effect of full parameter uncertainty reduces the economic value of volatility and correlation timing as investors will optimally choose less risky portfolios, and thus attain lower fees and lower Sharpe ratios. With full parameter uncertainty, the Sharpe ratios are 0.56 for MLR, 0.73 for CCC and 0.79 for DCC. Despite out-of-sample estimation and the uncertainty over a large number of parameters, there are still economic gains from volatility and correlation timing.

7 Conclusion

The empirical literature in financial economics has long determined that asset return volatilities and correlations vary over time. Therefore, accurate forecasts of volatilities and correlations are critical for an investor's optimal asset allocation. This has motivated a long line of research dedicated to developing tractable multivariate volatility models. This chapter provides a comprehensive evaluation of the economic value of these models for dynamic strategies that invest in the FX market. We focus on the portfolio choice of an investor who is uncertain about the parameter estimates and the model specification.

Our analysis shows that there is high economic value in timing both FX volatilities and correlations: the performance fee for volatility timing is about 4% per year and correlation timing adds a further 3% per year. This result is robust to reasonable transaction costs, which in FX trading are generally low. It is also robust to parameter uncertainty, which has little or no effect on the economic value of volatility

and correlation timing. We find that the model with the simplest structure in dynamic correlations and volatilities performs equally well as models with asymmetric correlations, richer correlation structure or alternative volatility specifications. Despite its simplicity, therefore, the DCC model is a powerful instrument in international asset allocation. In conclusion, both volatility and correlation timing matter to an international investor, and it pays to take dynamic FX volatilities correlations into consideration in asset allocation strategies.

Appendix A: Univariate Models for Volatility Timing

We estimate the multivariate correlation models under nine volatility specifications for each asset:

1. GARCH: Bollerslev (1986):

$$\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2.$$

2. Absolute Value GARCH (AVGARCH: Taylor, 1986):

$$\sigma_t = \omega + \alpha |u_{t-1}| + \beta \sigma_{t-1}.$$

3. Nonlinear GARCH (NARCH: Higgins and Bera, 1992):

$$\sigma_t^\tau = \omega + \alpha |u_{t-1}|^\tau + \beta \sigma_{t-1}^\tau.$$

4. Exponential GARCH (EGARCH: Nelson, 1991):

$$\ln\left(\sigma_t^2\right) = \omega + \alpha \varepsilon_{t-1} + \kappa\left(|\varepsilon_{t-1}| - E\left[|\varepsilon_t|\right]\right) + \beta \ln(\sigma_{t-1}^2); \quad \varepsilon_{t-1} = \frac{u_{t-1}}{\sigma_{t-1}}; \quad E\left[|\varepsilon_t|\right] = \sqrt{\frac{2}{\pi}}.$$

5. Threshold GARCH (ZARCH: Zakoian, 1994):

$$\sigma_t = \omega + \alpha\left(|u_{t-1}| - \kappa u_{t-1}\right) + \beta \sigma_{t-1}.$$

6. GJR-GARCH (Glosten, Jagannathan and Runkle, 1993):

$$\sigma_t^2 = \omega + \alpha\left(|u_{t-1}| - \kappa u_{t-1}\right)^2 + \beta \sigma_{t-1}^2.$$

7. Asymmetric Power GARCH (APGARCH: Ding, Engle and Granger, 1993):

$$\sigma_t^\tau = \omega + \alpha\left(|u_{t-1}| - \kappa u_{t-1}\right)^\tau + \beta \sigma_{t-1}^\tau.$$

8. Asymmetric GARCH (AGARCH: Engle, 1990):

$$\sigma_t^2 = \omega + \alpha\left(u_{t-1} + \kappa\right)^2 + \beta \sigma_{t-1}^2.$$

9. Nonlinear Asymmetric GARCH (NAGARCH: Engle and Ng, 1993):

$$\sigma_t^2 = \omega + \alpha\left(u_{t-1} + \kappa \sigma_t\right)^2 + \beta \sigma_{t-1}^2.$$

Appendix B: Parameter Uncertainty and the Predictive Density

We account for parameter uncertainty in Bayesian asset allocation by computing the predictive density of y_{t+1} as follows:

$$p\left(y_{t+1} \mid y_t\right) = \int_{\theta} p\left(y_{t+1}, \theta \mid y_t\right) d\theta = \int_{\theta} p\left(y_{t+1} \mid y_t, \theta\right) \pi\left(\theta\right) d\theta. \tag{11}$$

When the portfolio allocation problem is intertemporal, the solution should take into account the fact that the posterior distribution changes each period as the investor incorporates into her posterior beliefs information contained in new data realizations. This allows us to investigate jointly the effect of predictability and parameter uncertainty in dynamic asset allocation.[16] For instance, we can compare four different portfolio problems corresponding to four subjective data generating processes:

- *No parameter uncertainty, no predictability*: the investor uses the point estimates of the parameters and treats means, volatilities and correlations as constant over time. This is equivalent to using the MLR model in plug-in asset allocation.

- *No parameter uncertainty, predictability*: the investor uses the point estimates of the parameters but takes into account predictability in volatilities and correlations. This is equivalent to using the dynamic models in plug-in asset allocation.

- *Parameter uncertainty, no predictability*: the investor takes into account parameter uncertainty and treats means, volatilities and correlations as constant over time. This is equivalent to using the MLR model in Bayesian asset allocation.

- *Parameter uncertainty, predictability*: the investor accounts for both parameter uncertainty and predictability in volatilities and correlations. This is equivalent to using the dynamic models in Bayesian asset allocation.

[16] Due to the complexity and high dimension of the empirical models for dynamic correlation, the combination of parameter uncertainty and predictability we examine here is slightly different from the case of full Bayesian learning considered by, for example, Brandt *et al.* (2005). In our Bayesian asset allocation, given the full MCMC sample of the parameters, we update the posterior second moments in every time period. However, it is not feasible to also update the parameters of the time-varying second moments. The case of full Bayesian learning would be feasible if the moments were constant over time as in the case of the MLR model.

Table 1: Descriptive Statistics for Daily Exchange Rate Returns

The table summarizes the descriptive statistics for the daily percent exchange rate returns. The data range from January 1976 to December 2006 for a sample size of 8,069 daily observations.

	GBP	EUR	CHF	JPY
	Percent Returns			
$Mean$	-0.0004	0.0070	0.0095	0.0117
Std	0.620	0.653	0.736	0.664
$Skew$	-0.356	-0.005	0.083	0.609
$Kurt$	9.78	6.02	6.25	9.27
	Correlation Matrix			
GBP	1.000	0.557	0.495	0.336
EUR	0.557	1.000	0.819	0.491
CHF	0.495	0.819	1.000	0.543
JPY	0.336	0.491	0.543	1.000

Table 2: Posterior Means for the ADCC$_{diag}$−GARCH Model

The table presents the Bayesian MCMC estimates of the posterior means of the *ADCC$_{diag}$-GARCH* model applied on daily percent exchange rate returns from January 1976 to December 2006. The 95% highest posterior density (*HPD*) region for each parameter estimate (the shortest interval that contains 95% of the posterior distribution) is reported in parentheses, and the numerical Standard Error (*NSE*) in brackets.

	GBP	*EUR*	*CHF*	*JPY*
		Mean Parameters		
μ	0.0036 (−0.0080, 0.0151) [0.00009]	0.0114 (−0.0010, 0.0235) [0.00008]	0.0115 (−0.0027, 0.0261) [0.00011]	0.0092 (−0.0031, 0.0216) [0.00009]
		Volatility Parameters		
ω	0.0062 (0.0046, 0.0081) [0.00002]	0.0040 (0.0027,0.0056) [0.00001]	0.0049 (0.0030, 0.0073) [0.00002]	0.0046 (0.0030, 0.0066) [0.00002]
α	0.0682 (0.0579, 0.0790) [0.00009]	0.0636 (0.0543, 0.0732) [0.00009]	0.0545 (0.0460, 0.0634) [0.00008]	0.0647 (0.0550, 0.0752) [0.00010]
β	0.9175 (0.9043, 0.9296) [0.00012]	0.9298 (0.9194, 0.9394) [0.00010]	0.9383 (0.9276, 0.9480) [0.00010]	0.9283 (0.9157, 0.9395) [0.00012]
		Correlation Parameters		
γ	0.9809 (0.9782, 0.9833) [0.00004]	0.9780 (0.9765, 0.9797) [0.00003]	0.9793 (0.9775, 0.9810) [0.00003]	0.9871 (0.9854, 0.9886) [0.00003]
δ	0.1540 (0.1437, 0.1650) [0.00020]	0.1977 (0.1900, 0.2046) [0.00013]	0.1927 (0.1842, 0.2002) [0.00014]	0.1439 (0.1357, 0.1518) [0.00014]
π	0.1298 (0.1074, 0.1522) [0.00041]	0.0460 (0.0241, 0.0674) [0.00041]	0.0364 (0.0162, 0.0566) [0.00040]	0.0334 (0.0114, 0.0559) [0.00041]

Table 3: In-Sample Portfolio Performance

The table reports the in-sample economic value of selected currency strategies investing in the US dollar, British pound, Deutsche mark\euro, Swiss franc and Japanese yen. *MLR* is the benchmark strategy using the multivariate linear regression model, *CCC* is a dynamic strategy using the constant conditional correlation model, and *DCC* is a dynamic strategy using the dynamic conditional correlation model. The strategies build a portfolio by investing in five bonds from the US, UK, Germany, Switzerland and Japan and using the four exchange rate forecasts to convert the portfolio return in US dollars. The annualized percent mean, percent volatility and Sharpe ratio are denoted by μ_p, σ_p, and SR, respectively. Φ denotes the performance fee an investor with CRRA utility and a degree of relative risk aversion equal to 6 is willing to pay for switching from *MLR* to one of the dynamic strategies and is reported in annual basis points. The break-even transaction cost τ^{be} is defined as the minimum daily proportional cost that cancels out the utility advantage of a given strategy and is reported in daily basis points. The *Plug-In Method* uses mean, volatility and correlation forecasts without accounting for parameter uncertainty. *Parameter Uncertainty in Volatilities and Correlations* accounts for parameter uncertainty only in volatility and correlation forecasts. *Full Parameter Uncertainty* accounts for parameter uncertainty in all forecasts. The sample period covers daily data from January 1976 to December 2006.

Strategy	μ_p	σ_p	SR	Φ	τ^{be}
Plug-In Method					
MLR	18.8	11.2	1.08		
$CCC-GARCH$	22.7	11.8	1.36	386	10.5
$CCC-AVGARCH$	22.1	12.0	1.28	327	7.0
$CCC-NARCH$	22.4	11.9	1.32	356	8.4
$CCC-EGARCH$	22.9	12.1	1.34	404	8.5
$CCC-ZARCH$	22.6	12.1	1.32	377	8.1
$CCC-GJR-GARCH$	22.8	11.8	1.37	399	10.9
$CCC-APGARCH$	22.7	11.9	1.34	383	9.1
$CCC-AGARCH$	23.0	11.8	1.37	412	11.2
$CCC-NAGARCH$	23.0	11.8	1.38	416	11.3
$DCC-GARCH$	25.8	11.7	1.62	696	11.9
$DCC_{diag}-GARCH$	25.7	11.7	1.62	691	11.1
$ADCC-GARCH$	25.7	11.7	1.62	691	11.8
$ADCC_{diag}-GARCH$	25.8	11.7	1.62	694	11.1
Parameter Uncertainty in Volatilities and Correlations					
MLR	18.8	11.2	1.08		
$CCC-GARCH$	22.7	11.8	1.36	386	10.5
$DCC-GARCH$	25.7	11.7	1.62	692	11.8
Full Parameter Uncertainty					
MLR	18.5	10.8	1.08		
$CCC-GARCH$	22.3	11.4	1.36	371	9.9
$DCC-GARCH$	25.2	11.6	1.58	655	11.0

Table 4: In-Sample Portfolio Performance of Combined Forecasts

The table assesses the impact of model uncertainty on correlation timing by presenting the in-sample portfolio performance of combined forecasts. Expected utility is evaluated under the predictive density thus accounting for parameter uncertainty. *BMA* denotes *Bayesian Model Averaging* and *BMW Bayesian Model Winner*, which are applied on three universes of models: *VOL* is the universe of all GARCH-type univariate volatility specifications under the scalar symmetric *DCC* model; *CORR* is the universe of all multivariate correlation specifications (*CCC* and the four *DCC* specifications) with *GARCH* volatility; and *FULL* is the complete universe of all 46 model specifications (including the benchmark *MLR*). The annualized percent mean, percent volatility and Sharpe ratio are denoted by μ_p, σ_p, and SR, respectively. Φ denotes the performance fee an investor with CRRA utility and a degree of relative risk aversion equal to 6 is willing to pay for switching from *MLR* to one of the dynamic forecast combinations and is reported in annual basis points. The break-even transaction cost τ^{be} is defined as the minimum proportional cost that cancels out the utility advantage of a given strategy and is reported in daily basis points. The sample period covers daily data from January 1976 to December 2006.

Strategy	μ_p	σ_p	SR	Φ	τ^{be}
Bayesian Model Averaging					
$BMA-VOL$	24.2	10.6	1.65	568	7.9
$BMA-CORR$	25.1	11.6	1.58	653	10.6
$BMA-FULL$	24.0	10.0	1.73	554	7.2
Bayesian Model Winner					
$BMW-VOL$	24.8	10.4	1.73	631	3.4
$BMW-CORR$	25.3	11.5	1.60	664	8.2
$BMW-FULL$	23.8	9.4	1.81	539	2.2

Table 5: Out-of-Sample Portfolio Performance

The table shows the out-of-sample economic value of selected currency strategies investing in the US dollar, British pound, Deutsche mark\euro, Swiss franc and Japanese yen. *MLR* is the benchmark strategy using the multivariate linear regression model, *CCC* is a dynamic strategy using the constant conditional correlation model, and *DCC* is a dynamic strategy using the dynamic conditional correlation model. The strategies build a portfolio by investing in five bonds from the US, UK, Germany, Switzerland and Japan and using the four exchange rate forecasts to convert the portfolio return in US dollars. The annualized percent mean, percent volatility and Sharpe ratio are denoted by μ_p, σ_p, and SR, respectively. Φ denotes the performance fee an investor with CRRA utility and a degree of relative risk aversion equal to 6 is willing to pay for switching from *MLR* to one of the dynamic strategies and is reported in annual basis points. The break-even transaction cost τ^{be} is defined as the minimum daily proportional cost that cancels out the utility advantage of a given strategy and is reported in daily basis points. The *Plug-In Method* uses mean, volatility and correlation forecasts without accounting for parameter uncertainty. *Parameter Uncertainty in Volatilities and Correlations* accounts for parameter uncertainty only in volatility and correlation forecasts. *Full Parameter Uncertainty* accounts for parameter uncertainty in all forecasts. The sample period covers daily data from January 1976 to December 2006. The out-of-sample period uses daily observations from January 1986 to December 2006 and sequentially updates the parameter estimates month-by-month using a 10-year rolling window.

Strategy	μ_p	σ_p	SR	Φ	τ^{be}
		Plug-In Method			
MLR	11.2	10.7	0.57		
$CCC-GARCH$	16.0	12.8	0.86	457	7.8
$DCC-GARCH$	19.8	15.4	0.96	800	9.7
$DCC_{diag}-GARCH$	19.1	14.9	0.94	732	8.4
$ADCC-GARCH$	19.6	15.3	0.95	774	9.4
$ADCC_{diag}-GARCH$	18.7	14.9	0.91	693	7.6
		Parameter Uncertainty in Volatilities and Correlations			
MLR	11.2	10.7	0.57		
$CCC-GARCH$	16.0	12.8	0.86	460	6.6
$DCC-GARCH$	19.9	15.4	0.96	812	8.4
		Full Parameter Uncertainty			
MLR	11.4	11.2	0.56		
$CCC-GARCH$	16.2	15.2	0.73	428	4.2
$DCC-GARCH$	18.3	16.8	0.79	613	4.9

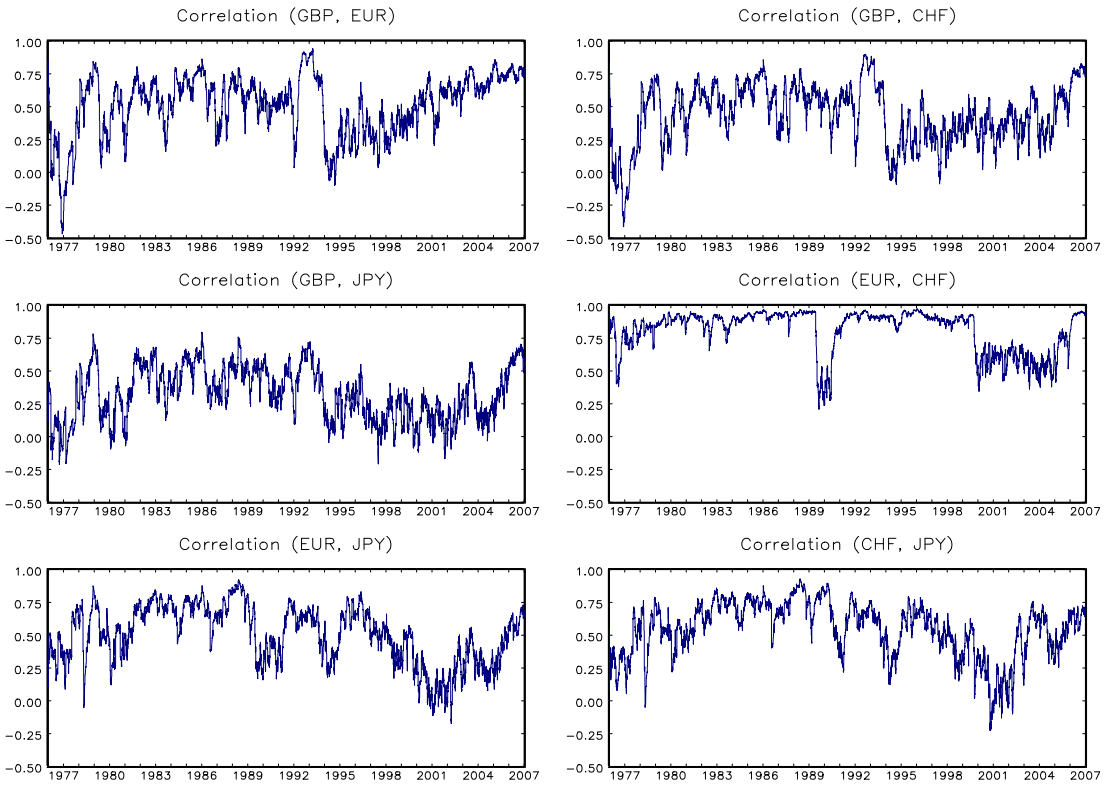

Figure 1. Daily Correlation Forecasts

The figure displays the in-sample daily correlation forecasts between four US dollar exchange rate return series using the simple (scalar symmetric) DCC-GARCH model. The in-sample period uses daily observations ranging from January 1976 to December 2006.

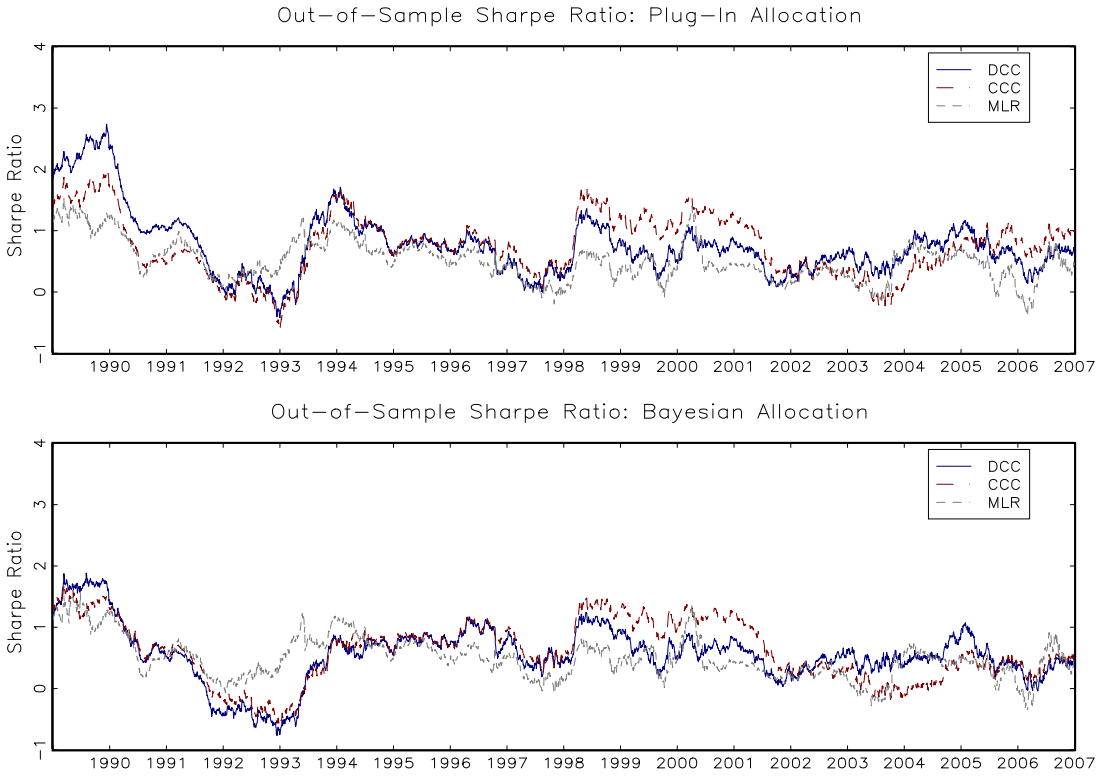

Figure 2. Out-of-Sample Sharpe Ratios

The figure displays the out-of-sample Sharpe ratios for three strategies. MLR is the benchmark strategy with static volatilities and correlations; CCC has dynamic GARCH volatilities but constant correlations; and DCC has dynamic (scalar symmetric) correlations and GARCH volatilities. The top panel is for plug-in allocation and the bottom panel for Bayesian allocation with full parameter uncertainty. All strategies are evaluated for a degree of relative risk aversion equal to 6. The figure displays the case of zero transaction costs. The out-of-sample period uses daily observations from January 1986 to December 2006 and sequentially updates the parameter estimates month-by-month using a 10-year rolling window. The Sharpe ratios are calculated using a 3-year rolling window and hence start in January 1989.

References

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

Ang, A., and G. Bekaert (2002). "International Asset Allocation with Regime Shifts," *Review of Financial Studies* **15**, 1137–1187.

Ang, A., and J. Chen (2002). "Asymmetric Correlations of Equity Portfolios," *Journal of Financial Economics* **63**, 443–494.

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

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

Bauwens, L., and S. Laurent (2005). "A New Class of Multivariate Skew Densities with Application to Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business and Economic Statistics* **23**, 346–354.

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

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.

Bollerslev, T., R.F. Engle, and J.M. Wooldridge (1988). "A Capital Asset Pricing Model With Time Varying Covariances," *Journal of Political Economy* **96**, 116–131.

Brandt, M.W. (2009). "Portfolio Choice Problems", in Y. Ait-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*, New York: Elsevier.

Brandt, M.W., A. Goyal, P. Santa-Clara, and J.R. Stroud (2005). "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning about Return Predictability," *Review of Financial Studies* **18**, 831–873.

Buraschi, A., P. Porchia, and F. Trojani, (2010). "Correlation Risk and Optimal Portfolio Choice," *Journal of Finance* **65**, 393–420.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011). "Do Peso Problems Explain the Returns to the Carry Trade?" *Review of Financial Studies* **24**, 853–891.

Cappiello, L., R.F. Engle, and K. Sheppard (2006). "Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns," *Journal of Financial Econometrics* **4**, 537–572.

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.

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

Della Corte, P., L. Sarno, and G. Sestieri (2011). "The Predictive Information Content of External Imbalances for Exchange Rate Returns: How Much Is It Worth?," *Review of Economics and Statistics* (forthcoming) doi:10.1162/REST_a_00157.

Della Corte, P., L. Sarno, and D.L. Thornton (2008). "The Expectation Hypothesis of the Term Structure of very Short-Term Rates: Statistical Tests and Economic Value," *Journal of Financial Economics* **89**, 158–174.

Della Corte, P., L. Sarno, and I. Tsiakas (2009). "An Economic Evaluation of Empirical Exchange Rate Models," *Review of Financial Studies* **22**, 3491–3530.

Della Corte, P., L. Sarno, and I. Tsiakas (2011). "Spot and Forward Volatility in Foreign Exchange," *Journal of Financial Economics* **100**, 496–513.

Della Corte, P., and I. Tsiakas (2012). "Statistical and Economic Methods for Evaluating Exchange Rate Predictability." in James, J., L. Sarno and I.W. Marsh (eds.) *Handbook of Exchange Rates*. London: Wiley (forthcoming).

Diebold, F.X., and M. Nerlove (1989). "The Dynamics of Exchange Rate Volatility: A Multivariate Latent Factor ARCH Model," *Journal of Applied Econometrics* **4**, 1–21.

Ding, Z., R.F. Engle, and C. Granger (1993). "A Long-Memory Property of Stock Market Returns and a New Model," *Journal of Empirical Finance* **1**, 83–106.

Driessen, J., P. Maenhout, and G. Vilkov (2009). "The Price of Correlation Risk: Evidence from Equity Options," *Journal of Finance* **64**, 1377–1406.

Engle, R.F. (1990). "Discussion: Stock Market Volatility and the Crash of '87," *Review of Financial Studies* **3**, 103–106.

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

Engle, R.F., and R. Colacito (2006). "Testing and Valuing Dynamic Correlations for Asset Allocation," *Journal of Business and Economic Statistics* **24**, 238–253.

Engle, R.F., and K.F. Kroner (1995). "Multivariate Simultaneous Generalized ARCH," *Econometric Theory* **11**, 122–150.

Engle, R.F., and V.K. Ng (1993). "Measuring and Testing the Impact of News on Volatility," *Journal of Finance* **48**, 1749–1778.

Engle, R.F., N. Shephard, and K. Sheppard (2008). "Fitting Vast Dimensional Time-Varying Covariance Models," Unpublished working paper, University of Oxford.

Engle, R.F., and K. Sheppard (2001). "Theoretical and Empirical Properties of Dynamic Conditional Correlation Multivariate GARCH," Unpublished working paper, University of California at San Diego.

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

Fleming, J., C. Kirby, and B. Ostdiek (2003). "The Economic Value of Volatility Timing using 'Realized' Volatility," *Journal of Financial Economics* **67**, 473–509.

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

Glosten, L.R., R. Jagannathan, and D.E. Runkle (1993). "On the Relationship Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance* **48**, 1779–1801.

Goetzmann, W.N., L. Li, and K.G. Rouwenhorst (2005). "Long Term Global Market Correlations," *Journal of Business* **78**, 1–38.

Gourieroux, C., J. Jasiak, and R. Sufana (2009). "The Wishart Autoregressive Process of Multivariate Stochastic Volatility," *Journal of Econometrics* **150**, 167–181.

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 Economic Studies* **61**, 247–264.

Higgins, M.L., and A.K. Bera (1992). "A Class of Nonlinear ARCH Models," *International Economic Review* **33**, 137–158.

Kan, R., and G. Zhou (2007). "Optimal Portfolio Choice with Parameter Uncertainty," *Journal of Financial and Quantitative Analysis* **42**, 621–656.

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

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

Ledoit, O., P. Santa-Clara, and M. Wolf (2003). "Flexible Multivariate GARCH Modeling With an Application to International Stock Markets," *Review of Economics and Statistics* **85**, 735–747.

Longin, F., and B. Solnik (1995). "Is the Correlation in International Equity Returns Constant: 1960-1990," *Journal of International Money and Finance* **14**, 3–26.

Longin, F., and B. Solnik (2001). "Extreme Correlation of International Equity Markets," *Journal of Finance* **56**, 649–675.

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.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2011), "Carry Trades and Global Foreign Exchange Volatility," *Journal of Finance* (forthcoming).

Nelson, D.B. (1991). "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica* **59**, 347–370.

Palandri, A. (2009) "Sequential Conditional Correlations: Inference and Evaluation," *Journal of Econometrics* **153**, 122–132.

Philipov, A., and M.E. Glickman (2006). "Multivariate Stochastic Volatility via Wishart Processes," *Journal of Business and Economic Statistics* **24**, 313–328.

Taylor, S.J. (1986). *Modeling Financial Time Series*. Chichester: J. Wiley & Sons Ltd.

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

Tsiakas, I. (2006). "Periodic Stochastic Volatility and Fat Tails," *Journal of Financial Econometrics* **4**, 90-135.

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.

Wright, J.H. (2008). "Bayesian Model Averaging and Exchange Rate Forecasting," *Journal of Econometrics* **146**, 329–341.

Zakoian, J.-M. (1994). "Threshold Heteroskedastic Models," *Journal of Economic Dynamics and Control* **18**, 931–955.

Zellner, A., and V.K. Chetty (1965). "Prediction and Decision Problems in Regression Models from the Bayesian Point of View," *Journal of the American Statistical Association* **60**, 608–616.