

Correlation Timing in Asset Allocation with Bayesian Learning*

Pasquale Della Corte^a Lucio Sarno^{a,b} Ilias Tsiakas^a

a: University of Warwick

b: Centre for Economic Policy Research (CEPR)

September 2008

Abstract

This paper investigates how dynamic correlations in exchange rate returns affect the optimal portfolio choice of risk-averse investors. We use a set of multivariate models based on the Dynamic Conditional Correlation (DCC) model and take a Bayesian approach in estimation and asset allocation. In assessing the economic value of correlation timing, our analysis accounts for parameter and model uncertainty as well as Bayesian learning. We design dynamic strategies for active management of currency portfolios, and find substantial economic value in timing correlations in addition to the economic gains from volatility timing. This result is robust to reasonable transaction costs as well as parameter and model uncertainty, asymmetric correlations, and alternative specifications for correlations and volatilities; it is not robust, however, to the portfolio rebalancing frequency as the large economic gains of correlation timing at the daily frequency dissipate almost entirely for weekly or monthly rebalancing.

Keywords: Asset Allocation; Correlation Timing; Volatility Timing; Bayesian Learning; Bayesian Model Averaging.

JEL Classification: F31; F37; G10; G11.

***Acknowledgements:** The authors are indebted for useful conversations or constructive comments to Torben Andersen, Federico Bandi, Luc Bauwens, Frank Diebold, Robert Engle, Rene Garcia, John Geweke, Eric Ghysels, Mico Loretan, John Maheu, Anna Pavlova, Tarun Ramadorai, 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. *Corresponding author:* Ilias Tsiakas, Finance Group, Warwick Business School, University of Warwick, Coventry CV4 7AL, UK. E-mail: ilias.tsiakas@wbs.ac.uk. *Other authors' contact details:* Pasquale Della Corte, pasquale.dellacorte@wbs.ac.uk; Lucio Sarno, lucio.sarno@wbs.ac.uk.

1 Introduction

Expected correlations between asset returns are a critical input in the optimal portfolio choice of risk-averse investors. Furthermore, there is considerable empirical evidence that return correlations change over time.¹ However, forecasting the dynamics of correlations 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 is the first paper to address an essential question that lies at the core of a long line of research in empirical finance: does correlation timing matter for the optimal asset allocation of a risk-averse investor and, if so, how? We thus provide a comprehensive analysis of the extent to which dynamic return correlations affect the optimal portfolio choice of risk-averse investors. Our analysis contributes to the literature on the economic value of volatility timing, which 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, 2008), 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.

Our empirical investigation begins by estimating a large set of multivariate specifications based on the DCC model. In doing so, we take a Bayesian approach in estimation and asset allocation. This allows us to make a further contribution to the literature by evaluating correlation timing in a way that recognizes the role of both parameter and model uncertainty. In the process, we also assess the impact of other important aspects of portfolio choice, such as transaction costs, rebalancing frequency, the choice of utility function, asymmetry in correlations, and richness of correlation dynamics.

The analysis of this paper focuses on the foreign exchange (FX) market by making use of 31 years of daily returns data from four major US dollar nominal spot exchange rates. As the largest financial market in the world, the FX market is geographically dispersed with a uniquely international dimension.³ More importantly, it is a natural market to study correlations as investors trade currencies

¹See, for example, Erb, Harvey and Viskanta (1994), Longin and Solnik (1995, 2001), Ball and Torous (2000), Campbell, Lettau, Malkiel and Xu (2001), Ang and Bekaert (2002), Goetzmann, Li and Rouwenhorst (2005), and Cappiello, Engle and Sheppard (2006).

²See, 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), Fiorentini, Sentana and Shephard (2004), Chib, Nardari and Shephard (2006), Philipov and Glickman (2006), and Gouriéroux, Jasiak and Sufana (2008).

³The FX market has an average daily volume of transactions exceeding US \$3 trillion (see Bank of International Settlements, 2007).

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 in the US economy will move the US dollar in the same direction relative to all other currencies, thus generating positive correlation in all dollar exchange rates. Therefore, the consequence of pricing relative to a numeraire is the tendency of FX correlations to be positive. More generally, 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.

There are also practical reasons why focusing on the FX market makes our analysis more compelling. In particular, we assess the relative economic value of correlation and volatility timing without modeling exchange rate returns as a function of state variables. 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).⁴ Furthermore, 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 taking both long and short positions. Finally, in recent years investors can directly trade on FX correlations using an increasingly popular contract called the correlation swap.⁵

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 correlation models, there is little work in formally assessing the economic value of correlation timing.⁶ A purely statistical analysis of correlation timing is not particularly informative to an investor as it falls short of measuring whether there are tangible economic gains from implementing dynamic correlations in active portfolio management. This motivates our asset allocation approach, which extends previous studies of volatility timing in asset returns by West, Edison and Cho (1993), Fleming, Kirby and Ostdiek (2001), Markering and Verbeek (2004), Han (2006), and Della Corte, Sarno and Tsiakas (2008). We evaluate the dynamic allocation strategies using a constant relative risk aversion (CRRA) utility function,

⁴On the one hand, Meese and Rogoff (1983) find that exchange rate returns are not predictable since models that condition on economic fundamentals cannot outperform a naive random walk model. On the other hand, numerous empirical studies consistently reject the uncovered interest parity condition suggesting the lack of a meaningful economic relation between interest rate differentials and future exchange rate returns (e.g., Engel, 1996). These findings justify the widespread use of the random walk model as the basis of carry trade FX strategies that borrow in low interest rate currencies and lend in high interest rate currencies.

⁵In exchange for a fixed payment, a correlation swap pays the average realized correlation on a basket of securities. A dispersion trade, widely used by hedge funds and proprietary traders, can also attain similar exposure to correlation risk by taking a position in an index (option) and the opposite position in all components of the index (option).

⁶Engle and Sheppard (2001) examine the return volatility of global minimum variance portfolios implied by the DCC model, but do not explicitly measure the utility cost of constant conditional correlations. Engle and Colacito (2006) introduce a set of formal statistical tests for comparing covariance estimators in the context of minimum volatility allocation strategies. However, their study does not provide a utility evaluation of dynamically rebalanced portfolios from the perspective of a risk-averse investor.

thus departing from the standard quadratic utility mean-variance setting adopted by the majority of prior research in volatility timing. Ultimately, we 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 correlation and volatility in the presence of parameter and model uncertainty.

We assess the economic value of correlation timing in a Bayesian framework, which allows us to explicitly account for the fact that the true extent of correlation timing in FX returns is uncertain. Since the correlation forecasts are not known with complete precision, 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. This holds because the predictive distribution is determined by historical data and the prior, but does not depend on the parameter estimates (e.g., Kandel and Stambaugh, 1996; Barberis, 2000; Kan and Zhou, 2007; and Skoulakis, 2007). We can thus examine the effect of parameter uncertainty on asset allocation by comparing, on the one hand, the “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. Finally, the dynamic nature of the portfolio choice problem implies that the predictive distribution changes each period as the investor incorporates information contained in the new data realization into posterior beliefs. This feature of the problem introduces the notion of Bayesian learning, which may affect the investor’s portfolio holdings.

In line with the Bayesian approach of Avramov (2002), Cremers (2002), and Della Corte, Sarno and Tsiakas (2008), we also evaluate the impact of model risk on correlation timing by exploring whether portfolio performance improves when combining the forecasts arising from the large set of models we estimate.⁷ In fact, including the benchmark static covariance model, we estimate a total of 46 model specifications, and then use them in optimally implementing Bayesian Model Averaging (BMA). The BMA strategy accounts directly for uncertainty in model selection as it weighs all conditional volatility and correlation forecasts by the posterior probability of each model. The posterior probability is based on the marginal likelihood, and hence, it also accounts for parameter uncertainty, while imposing a penalty for lack of parsimony (higher dimension).

To preview our key results, the performance of the dynamic allocation strategies suggests that there is high economic value in timing the correlations of exchange rate returns. We find that an international investor facing FX risk will pay a high performance fee for correlation timing over

⁷We estimate the parameters of the DCC model by designing a new 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.

and above the case of volatility timing with constant correlations. This finding is robust to reasonable transaction costs as well as parameter uncertainty, model risk, asymmetric correlations, and alternative specifications for correlations and volatilities. However, this result holds only for daily portfolio rebalancing as the economic value of correlation timing dissipates almost entirely for weekly or monthly rebalancing. Furthermore, the simplest DCC specification examined here captures almost all gains from timing correlations and volatilities in the context of asset allocation across currencies. This is an important new finding that answers an essential question in the empirical finance literature that over the years has developed increasingly sophisticated multivariate volatility models. In the practice of currency management, therefore, forecasting correlations using the simplest DCC model will significantly enhance the performance of optimally designed portfolios when rebalancing at daily frequency.

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

2 Dynamic Models for Correlation Timing

We model the dynamics of correlations and volatilities of exchange rate returns using a set of models 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 relative economic value of correlation and volatility 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 for Correlation Timing

We begin our formal description of the models by letting $y_t = (y_{1,t}, \dots, y_{N,t})'$ denote the returns of N assets at time t :

$$y_t = \mu_t + \Sigma_t^{1/2} \varepsilon_t, \quad (1)$$

where $\mu_t = (\mu_{1,t}, \dots, \mu_{N,t})'$ is the vector of conditional means, Σ_t is the conditional covariance matrix, and $\varepsilon_t = (\varepsilon_{1,t}, \dots, \varepsilon_{N,t})'$ is a vector of standard normal disturbances.⁸

As the primary focus of our study is the effect of dynamic correlation and volatility on asset allocation, our analysis assumes a constant conditional mean $\mu_t = \mu$ for all models. This is equivalent to specifying a random walk model for the spot log-exchange rate, which in turn is consistent with most of the empirical FX literature (Meese and Rogoff, 1983; Engel, Mark and West, 2007).

2.1.1 The Static Benchmark

The *Multivariate Linear Regression (MLR)* model sets $\Sigma_t = \Sigma$, and therefore, it is the static benchmark, which is equivalent to a multivariate random walk of log-exchange rates with a constant covariance matrix.

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:

$$\Sigma_t = D_t \bar{R} D_t, \quad (2)$$

$$D_t = \text{diag} \{ \sigma_{1,t}, \dots, \sigma_{N,t} \}, \quad (3)$$

where D_t is the $N \times N$ diagonal matrix of dynamic conditional volatilities, and \bar{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 correlations in addition to dynamic volatilities by decomposing the conditional covariance matrix as follows:

$$\Sigma_t = D_t R_t D_t, \quad (4)$$

$$D_t = \text{diag} \{ \sigma_{1,t}, \dots, \sigma_{N,t} \}, \quad (5)$$

⁸For recent studies which generalize the conditional normality assumption of the DCC model see Bauwens and Laurent (2005) and Jondeau and Rockinger (2008).

$$R_t = (\text{diag}\{Q_t\})^{-1/2} Q_t (\text{diag}\{Q_t\})^{-1/2}, \quad (6)$$

$$Q_t = (\bar{R} - \Gamma' \bar{R} \Gamma - \Delta' \bar{R} \Delta) + \Gamma' z_{t-1} z'_{t-1} \Gamma + \Delta' Q_{t-1} \Delta, \quad (7)$$

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, \bar{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 $u_t = y_t - \mu_t$.

The simplest version of the *DCC* model reduces $\Gamma = \gamma$ and $\Delta = \delta$, where $\{\gamma, \delta\}$ are scalars, which are the same for all assets $i \leq N$. In other words, this “scalar” *DCC* model 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: $\Gamma = \text{diag}\{\gamma_1, \dots, \gamma_N\}$ and $\Delta = \text{diag}\{\delta_1, \dots, \delta_N\}$. The diagonal *DCC* model allows for distinct dynamics in each correlation process but requires estimation of more parameters. We estimate both scalar (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:

$$Q_t = (\bar{R} - \Gamma' \bar{R} \Gamma - \Delta' \bar{R} \Delta - \Pi' \bar{P} \Pi) + \Gamma' z_{t-1} z'_{t-1} \Gamma + \Delta' Q_{t-1} \Delta + \Pi' p_{t-1} p'_{t-1} \Pi, \quad (8)$$

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; $\bar{P} = E[p_t p'_t]$. 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 *DCC* model, 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; Das and Uppal, 2004). When asset returns tend to be 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.⁹ In this context, the *ADCC* model allows us to determine the

⁹Recent 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., Buraschi, Porschia and Trojani, 2008; Driessen, Maenhout and Vilkov, 2008; Krishnan, Petkova and Ritchken, 2008).

possible impact of asymmetric correlations on asset allocation in the FX market.

2.2 The Set of Univariate Models for Volatility Timing

We evaluate the performance of a number of alternative univariate specifications for the conditional variance, including some of the most popular models in the literature. The univariate volatility models we consider are the same as in Cappiello, Engle and Sheppard (2006) and include the following: 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). For details on the full specification of these models, see Appendix A.

2.3 Pairwise Model Comparisons

In addition to the static benchmark MLR , the set of models comprises five competing dynamic specifications $\{CCC, DCC, DCC_{diag}, ADCC$ and $ADCC_{diag}\}$ under each of the nine univariate volatility specifications listed above. In total, therefore, we estimate 46 model specifications.¹⁰

The principal objective of our analysis is to provide an economic evaluation of the models described above in the context of dynamic asset allocation strategies. We assess the economic value of volatility timing simply by comparing the CCC model to the static MLR . More importantly, we measure the additional economic gains from correlation timing by comparing the DCC to the CCC model. The models also allow us to 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.

2.4 Estimation and Forecasting

We perform Bayesian estimation of the parameters of all model specifications. The critical advantage of the Bayesian methodology is that it provides a unified framework for estimation, forecasting and model selection. In this application, it is particularly suitable for capturing the effect of parameter and model uncertainty on asset allocation. A Bayesian approach also avoids the complication on testing the null of constant correlation against the alternative of dynamic correlation that arises due to lack of identification of the correlation decay parameters (see Engle and Sheppard, 2001).

¹⁰Since our analysis focuses on exchange rates, we have also attempted to model the dynamic covariance between two dollar exchange rate returns using the dynamic variance of the cross-rate. This is based on triangular no-arbitrage and follows Andersen, Bollerslev, Diebold and Labys (2003) and Brandt and Diebold (2006). Despite the simplicity of reducing estimation of multivariate covariances to a set of univariate volatility processes, this model cannot guarantee that the dynamic covariance matrix is positive definite, and that correlations are in the $[-1,1]$ range. Hence we have excluded the triangular no-arbitrage model from our analysis.

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 as well as Bayesian learning.

In this paper, we innovate by designing a new Markov Chain Monte Carlo (MCMC) algorithm for Bayesian estimation of the *DCC* model. The algorithm draws insights from the Bayesian univariate *GARCH* algorithm of Della Corte, Sarno and Tsiakas (2008), and from the Bayesian stochastic volatility algorithm of Kim, Shephard, and Chib (1998) and Chib, Nardari, and Shephard (2002, 2006). 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 described in Section 4.

The MCMC algorithm for three representative multivariate models (*MLR*, *CCC*, *ADCC_{diag}*) is summarized in Appendix B. All other specifications are special cases of one of these models. Each algorithm produces estimates of the posterior means of $\theta = \{\mu, \theta_1, \theta_2\}$, where $\mu = \{\mu_i\}$ is the unconditional mean of each series $i \leq N$, 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.

3 The Economic Value of Correlation Timing

3.1 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 correlations and volatilities. 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. This investor problem is formally defined as follows (see, for example, Brandt, 2009):

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

where $R_{p,s+1} = x'_s (R_{s+1} - R^f) + R^f$ 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 . In this intertemporal portfolio choice problem, at date t the investor optimally chooses the portfolio weights x_t conditional on having wealth 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 function $V_t(W_t, Z_t)$ is called the value function and 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 $\{x_s^*\}_{s=t}^{T-1}$.

The dynamic nature of this portfolio choice is best understood by expressing the multi-period problem as a recursive single-period problem:

$$\begin{aligned} V_t(W_t, z_t) &= \max_{\{x_s\}_{s=t}^{T-1}} E_t[U(W_T)] \\ &= \max_{x_t} E_t \left\{ \max_{\{x_s\}_{s=t+1}^{T-1}} E_{t+1}[U(W_T)] \right\} \\ &= \max_{x_t} E_t \left\{ V_{t+1} \left[W_t \left(x'_t (R_{t+1} - R^f) + R^f \right), Z_{t+1} \right] \right\}. \end{aligned} \quad (10)$$

The second equality follows from the law of iterated expectations and the third equality uses the definition of the value function as well as the budget constraint. The recursive Equation (10) is the Bellman equation.

Consider the case of CRRA utility, which is defined as follows:

$$U(W_T) = \frac{W_T^{1-\lambda}}{1-\lambda},$$

where λ denotes the coefficient of relative risk aversion (RRA). We can show that for CRRA utility the Bellman equation simplifies to:

$$V_t(W_t, Z_t) = \max_{x_t} E_t \left[\underbrace{\frac{[W_t (x'_t (R_{t+1} - R^f) + R^f)]^{1-\lambda}}{1-\lambda}}_{U(W_{t+1})} \underbrace{\max_{\{x_s\}_{s=t+1}^{T-1} E_{t+1} \left\{ \prod_{s=t+1}^{T-1} (x'_s (R_{s+1} - R^f) + R^f) \right\}^{1-\lambda}}}_{\psi_{t+1}(Z_{t+1})} \right]. \quad (11)$$

The CRRA utility function is homogeneous of degree $1 - \lambda$ in wealth, which implies that the solution is invariant to wealth. Hence, without loss of generality, we can set $W_t = 1$. It follows that the value function depends only on the state variables as follows:

$$\frac{1}{1-\lambda} \psi_t(Z_t) = \max_{x_t} E_t \left[U \left(x'_t (R_{t+1} - R^f) + R^f \right) \psi_{t+1}(Z_{t+1}) \right]. \quad (12)$$

The corresponding first order conditions are:

$$E_t \left[\partial U \left(x'_t (R_{t+1} - R^f) + R^f \right) \psi_{t+1}(Z_{t+1}) (R_{t+1} - R^f) \right] = 0. \quad (13)$$

3.2 Optimal Portfolio Choice

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 for solving discrete-time portfolio choice problems, which is general in that it 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 correlation and volatility 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. The Brandt *et al.* (2005) solution to the dynamic allocation problem is discussed in Appendix C.¹¹

Our analysis employs the set of models for conditional correlations and volatilities discussed in Section 2. Since we are primarily interested in the relative economic value of volatility and correlation timing, all models assume a constant mean for exchange rate returns, which is distinct for each asset. By design, in this setting the optimal weights will vary across models only to the extent that forecasts of the conditional correlations and volatilities will vary, which is precisely what the empirical models provide. The benchmark against which we compare the model specifications is the static-covariance model (*MLR*). In short, our objective is to determine whether in FX markets: (i) there is economic value in conditioning on dynamic volatility, (ii) there is additional value in conditioning on dynamic correlations, (iii) if so, which volatility and correlation specification works best; (iv) whether parameter uncertainty and Bayesian learning have a significant impact on the optimal allocation; and finally, (v) whether we can extract further economic gains by accounting for model risk and constructing combined forecasts using Bayesian Model Averaging.

3.2.1 Performance Measures

At any point in time, one set of estimates of the conditional correlations and volatilities is better than a second set if investment decisions based on the first set lead to higher utility. Alternatively, the optimal model requires less wealth to yield a given level of utility than a suboptimal model. Following Fleming, Kirby and Ostdiek (2001), we measure the economic value of the 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 *MLR* model yields the same average utility as holding the *DCC* optimal portfolio that is subject to daily expenses Φ , expressed as a fraction of wealth invested in the portfolio. Since the investor would be indifferent between these two

¹¹Note that there is one important difference in our approach and the solution of Brandt *et al.* (2005). Since we directly estimate the conditional moments, we do not need to run regressions across simulated sample paths, and hence, our solution is slightly simpler.

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 better dynamic correlation and volatility forecasts. To estimate the performance fee, we find the value of Φ that satisfies:

$$\sum_{t=0}^{T-1} E_t [U(R_{p,t+1}^* - \Phi)] = \sum_{t=0}^{T-1} E_t [U(R_{p,t+1})], \quad (14)$$

where $R_{p,t+1}^*$ is the gross portfolio return constructed using the expected return, volatility and correlation forecasts from a competing dynamic model (such as the *DCC* model), and $R_{p,t+1}$ is the gross portfolio return implied by the benchmark *MLR* model.

The most commonly used measure of economic value is the Sharpe ratio. However, as suggested by Marquering and Verbeek (2004) and Han (2006), the realized Sharpe ratio can be misleading in evaluating the performance of dynamic strategies because it is computed using the sample standard deviation of realized portfolio returns, which overestimates the volatility of the portfolio returns of a dynamic strategy. Furthermore, the Sharpe ratio is based on a notion of risk that ignores higher moments and cannot quantify the exact economic gains of the dynamic strategies over the static covariance strategy in the direct way of performance fees. Therefore, our economic analysis of short-horizon exchange rate predictability focuses primarily on performance fees.

3.3 The Dynamic Strategies

We design an international strategy that involves trading the five major currencies (four exchange rates) under investigation. Consider a US investor who builds a portfolio by allocating her wealth between five bonds: one domestic (US), and four foreign bonds (UK, Germany, Switzerland and Japan). At the beginning of each period, the five bonds yield a riskless return in local currency. Hence the only risk the US investor is exposed to is FX risk.¹² Every period the investor takes two steps. First, she uses each model to forecast the one-day ahead conditional correlations and volatilities 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 correlation and volatility specification affects the performance of an allocation strategy in an economically meaningful way. In assessing the relative economic value of correlation and volatility timing, we consider three portfolio rebalancing frequencies: daily, weekly and monthly. The yield of the riskless bonds is proxied by daily eurodeposit rates. Following Marquering and Verbeek (2004) and Han (2006) we set $\lambda = 6$, which produces portfolios with reasonable expected return and volatility. We report the estimates

¹²Strictly speaking, the return of investing in a risky asset is equal to the certain yield of the foreign bond plus the uncertain exchange rate return (y_t). The risk-free rate is equal to the yield of the US bond.

of Φ as annualized fees in basis points.

3.4 Transaction Costs

The impact of transaction costs is an essential consideration in assessing the profitability of trading strategies. This is especially true here because the trading strategy based on the *MLR* benchmark is static, whereas the competing dynamic models generate dynamic strategies.¹³ Therefore, we calculate the performance fees of the dynamic strategies relative to the static benchmark for the case when the proportional transaction cost to the portfolio return is equal to 0, 1 and 2 basis points per trading period. In foreign exchange trading, this is a realistic range for the level of transaction costs as it roughly corresponds to 1 or 2 pips per currency unit.¹⁴ We follow the simple approximation of Marquering and Verbeek (2004) by deducting the proportional transaction cost from the portfolio return ex post. This ignores the fact that dynamic portfolios are no longer optimal in the presence of transaction costs but maintains simplicity and tractability in our analysis.

We can avoid these concerns by calculating the break-even proportional transaction cost, τ^{BE} , that renders investors indifferent between two strategies (e.g., Han, 2006; Della Corte, Sarno and Tsiakas, 2008). In comparing a dynamic strategy with the static 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 basis points at the rebalancing frequency. For example, for daily portfolio rebalancing we report τ^{BE} in daily basis points.

4 Parameter Uncertainty and Bayesian Learning

Asset allocation theories typically assume that investors make optimal decisions with full knowledge of the true parameters of the model. In practice, however, model parameters are not fixed at a known value and instead have to be estimated. If there is estimation error, the resulting allocation will be suboptimal. This gives rise to estimation risk in all empirical applications of the plug-in method that replaces the true parameter values by their estimates. In contrast, the Bayesian approach to asset allocation provides a general framework that integrates estimation risk into the analysis. Bayesian analysis 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 return predictability but is uncertain about the

¹³The *MLR* model assumes a constant covariance matrix, and hence the optimal weights fluctuate only to the extent that domestic and foreign interest rates vary over time.

¹⁴The 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 2.0000, 1 pip would raise it to 2.0001 and this would roughly correspond to 1/2 basis point proportional cost.

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 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), Kan and Zhou (2007), and Skoulakis (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 because it implies 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. Following Geweke (1989a), the predictive density of y_{t+1} is given by:

$$p(y_{t+1} | y_t) = \int_{\theta} p(y_{t+1}, \theta | y_t) d\theta = \int_{\theta} p(y_{t+1} | y_t, \theta) \pi(\theta) d\theta. \quad (15)$$

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 information contained in the new data realization into posterior beliefs. This feature of the problem introduces the notion of Bayesian updating or learning. Our dynamic asset allocation framework for assessing the relative economic value of correlation and volatility timing separates the effect of parameter uncertainty from that of learning. For instance, we can compare four different portfolio problems corresponding to four different subjective data generating processes:

- *No parameter uncertainty, no predictability*: the investor takes the parameters as constant at the point estimates and also treats means, volatilities and correlations as constant over time. This is equivalent to using the *MLR* model in the plug-in asset allocation.
- *No parameter uncertainty, predictability*: the investor takes the parameters as constant at the point estimates but allows for dynamics in volatilities and correlations. This is equivalent to using the dynamic models in the plug-in asset allocation.
- *Parameter uncertainty, no predictability*: the investor takes into account parameter uncertainty but does not learn from new data realizations (i.e., no predictability), as in Barberis (2000). This involves using the *MLR* model in Bayesian asset allocation.
- *Parameter uncertainty and predictability (Bayesian learning)*: the investor accounts for parameter uncertainty while learning about the return-generating process through the Bayesian predictive density due to predictability in correlations and volatilities. This is equivalent to using the dynamic models in Bayesian asset allocation.

In short, our plan for understanding the effect of parameter uncertainty and Bayesian learning on correlation timing is to compare, on the one hand, the allocation of an investor who uses the predictive distribution in forecasting 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. As a result, this is the first paper to explicitly incorporate learning about dynamic volatilities and correlations in a Bayesian multivariate asset allocation framework.

5 Model Risk and Combined Forecasts

Model risk arises from the uncertainty over selecting a model specification. Consistent with our Bayesian approach, a natural criterion for resolving this uncertainty is the posterior probability of each model, which has three important advantages: (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.

Consider a set of K models M_1, \dots, M_K . We form a prior belief $\pi(M_i)$ on the probability that the i th model is the true model, observe the returns data $\{y_t\}$, and then update our belief that the i th model is true by computing the posterior probability of each model as follows:

$$p(M_i | y) = \frac{p(y | M_i) \pi(M_i)}{\sum_{j=1}^K p(y | M_j) \pi(M_j)}, \quad (16)$$

where $p(y | M_i)$ is the marginal likelihood of the i th model defined as:

$$p(y | M_i) = \int_{\theta} p(y, \theta | M_i) d\theta = \int_{\theta} p(y | \theta, M_i) \pi(\theta | M_i) d\theta. \quad (17)$$

In Equation (16) above we set our prior belief to be that all models are equally likely, i.e. $\pi(M_i) = \frac{1}{K}$. 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.¹⁵

Furthermore, the marginal likelihood is simply the normalizing constant of the posterior density and (suppressing the model index for simplicity) it can be written as:

$$p(y) = \frac{f(y | \theta) \pi(\theta)}{\pi(\theta | y)}, \quad (18)$$

where $f(y | \theta)$ is the likelihood, $\pi(\theta)$ the prior density of the parameter vector θ , $\pi(\theta | y)$ the posterior density, and θ is evaluated at the posterior mean.¹⁶

¹⁵ *Occam’s Razor* is the principle of parsimony, which states that among two competing theories that make exactly the same prediction, the simpler one is best.

¹⁶ Since θ is drawn in the context of MCMC sampling, the posterior density $\pi(\theta | y)$ is computed using the technique of reduced conditional MCMC runs of Chib (1995). For the parameters sampled in the MCMC chain by implementing a Metropolis-Hastings algorithm, the posterior density is computed as in Chib and Jeliazkov (2001).

5.1 Combined Forecasts

Since we do not know which one of the competing models is true, we also assess the performance of combined forecasts proposed by the seminal work of Bates and Granger (1969). We design two strategies based on a combination of forecasts for the conditional volatility and correlation of exchange rate returns: the Bayesian Model Average (*BMA*) strategy and the Bayesian Model Winner (*BMW*) strategy.¹⁷ 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, Φ , for the *BMA* one-step ahead forecasts of the conditional volatilities and correlations and then compare them to the *MLR* benchmark. We perform this exercise for three universes of models: (i) *VOL* is the universe of all *GARCH*-type univariate volatility specifications under the scalar symmetric *DCC* model (the simplest *DCC* model we consider); (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*. Consequently, our analysis of correlation timing further contributes to the empirical finance literature by incorporating an economic view of both parameter and model uncertainty.

5.1.1 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 straightforward to implement once we have the output from the MCMC simulations. 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 defined in Equation (16). Note that the *BMA* weights vary not only across models but also across time as does the predictive density (and hence marginal likelihood) of each model. Finally, it is crucial to emphasize that we evaluate the *BMA* strategy ex-ante.¹⁸

5.1.2 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 empirical model that has the highest predictive density in that period. In other words, every period the *BMW* strategy only uses the forecasts of the “winner” model in terms of predictive density, and hence discards the forecasts of the other models. Clearly,

¹⁷See Timmermann (2006) for a review of forecast combinations.

¹⁸In practice, we do this by lagging the posterior probability of each model; see Della Corte, Sarno and Tsiakas (2008) for further discussion.

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 for 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 UK pound sterling (USD/GBP), the Deutsche mark/euro (USD/EUR), the Swiss franc (USD/CHF), and the Japanese yen (USD/JPY). After the introduction of the euro in January 1999, we use the official euro-Deutsche mark conversion rate to obtain the USD/EUR series. The weekly and monthly data we use for portfolio rebalancing at these frequencies are generated by sampling the Tuesday FX rates for the weekly data and the end-of-month rates for the monthly data.

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 USD/GBP to 0.0117 (or 2.9% per annum) for USD/JPY. The daily standard deviations revolve between 0.620 for USD/GBP (or 9.8% per annum) to 0.736 for USD/CHF (or 11.7% per annum). Skewness is negative for two of the four FX returns, while kurtosis ranges from 6.02 for USD/EUR to 9.78 for USD/GBP. Finally, Panel B of Table 1 shows the average return cross-correlations, which are strongly positive ranging between 0.336 and 0.819 .

6.2 Bayesian Estimation of Correlation Models

We begin our evaluation of correlation timing in exchange rate returns by performing Bayesian estimation of the parameters of all models set out in Section 2. In addition to the static *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 for daily, weekly and monthly returns. This is the most general multivariate correlation 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 correlations may be as predictable as volatilities. The persistence parameters are similar across assets, while their value declines slightly as we move from daily to weekly, and then to monthly frequency.¹⁹ Furthermore, the parameters $\{\pi_i\}$ indicating asymmetry in dynamic correlations are

¹⁹For instance, in the case of the symmetric scalar *DCC* model with *GARCH* volatility, the persistence in volatilities is captured by $\alpha + \beta$, while the persistence in correlations is $\gamma^2 + \delta^2$. The average $\alpha + \beta$ term across the four exchange

small for all daily exchange rates, with the exception of the UK pound. Note that in our Bayesian framework, we assess statistical significance by computing the Numerical Standard Error (NSE) as in Geweke (1989b). All parameter estimates exhibit very low NSE values, and therefore, a high degree of statistical significance.²⁰

With the parameter estimates at hand, we generate the correlation and volatility forecasts used in the asset allocation. Figure 1 illustrates the daily volatility forecasts from the simple $DCC-GARCH$ model, whereas Figure 2 shows the daily correlation forecasts from the same model. As expected, conditional correlations appear to fluctuate over time as much as conditional volatilities.

6.3 Evaluating Correlation Timing in Asset Allocation

We assess the economic value of exchange rate 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 ex ante (i.e., using only lagged information) but in-sample for three main reasons: (i) we estimate multivariate models for the returns of four exchange rates over a long daily sample spanning 31 years and for three frequencies (daily, weekly, monthly); (ii) Bayesian MCMC estimation is computationally demanding; and, more importantly, (iii) the universe of models we consider is very large comprising 46 model specifications. Consequently, extending our analysis out of sample (i.e., re-estimating the models period-by-period) is hardly feasible. Relying solely on in-sample results does not capture the forecasting power a practitioner might have had in real time. However, Inoue and Kilian (2004, 2006) provide a formal econometric analysis, which shows that both in-sample and out-of-sample tests of predictability are equally reliable under the null hypothesis of no predictability. In fact, their analysis concludes that in most cases in-sample tests have higher power, and therefore, tend to be more credible than out-of-sample tests. Using similar arguments, Cochrane (2008) provides evidence that favours in-sample tests in the context of stock return predictability.

Our economic evaluation focuses on the performance fee, Φ , a US investor is willing to pay for switching from one international strategy with FX risk to another. The fees are reported in Table 3 for daily rebalancing, Table 4 for weekly rebalancing, and Table 5 for monthly rebalancing. The tables display the economic value of each volatility and correlation specification relative to the benchmark MLR model. In our calculations, we follow Marquering and Verbeek (2004) and Han (2006) in setting a degree of relative risk aversion $\lambda = 6$, which generates portfolio returns with

rate returns falls slightly from 0.99 for daily returns to 0.97 for weekly returns and stays constant at 0.97 for monthly returns. The average $\gamma^2 + \delta^2$ term falls much faster from 0.99 for daily returns to 0.95 for weekly returns and to 0.84 for monthly returns.

²⁰The Numerical Standard Error is defined as: $NSE = \left\{ \frac{1}{I} \left[\hat{\psi}_0 + 2 \sum_{j=1}^{B_I} Ker(z) \hat{\psi}_j \right] \right\}^{\frac{1}{2}}$, where $I = 5,000$ is the number of iterations (beyond the initial burn-in of 1,000 iterations), $j = 1, \dots, B_I = 500$ lags is the set bandwidth, $z = \frac{j}{B_I}$, and $\hat{\psi}_j$ is the sample autocovariance of the MCMC draws for each estimated parameter cut according to the Parzen kernel $Ker(z)$.

reasonable mean and volatility. In Panel A of the Tables we report the results for CRRA utility under the classical plug-in method. In Panel B, we allow for Bayesian learning only in volatilities and correlations, but not in the mean. Panel C presents the results with full Bayesian learning as we evaluate the CRRA utility under the Bayesian predictive density, thus accounting for parameter uncertainty in the mean, volatilities and correlations. Panel A reports the results for an extensive set of models, whereas Panels B and C focus on a smaller number of selected models. The primary focus of our discussion below is the case of daily rebalancing.

In line with the evidence of, among others, West, Edison and Cho (1993), Fleming, Kirby and Ostdiek (2001), and Della Corte, Sarno and Tsiakas (2008), we find that there is substantial economic value associated with volatility timing. More importantly, however, we show that there is high economic value specifically due to timing dynamic FX correlations over and above the economic value of volatility timing. This result is a novel finding in this literature as we establish that for daily FX returns, the dynamics of covariances are not exclusively driven by dynamic volatilities, but also by dynamic correlations. Therefore, this is the first study to establish that correlation timing does matter to an international investor facing FX risk. We illustrate this finding by the following results.

In terms of portfolio performance, switching from the static *MLR* to the *CCC – GARCH* model gives a staggering fee of 386 annual basis points (*bps*). Then, switching from *MLR* to the *DCC – GARCH* model raises the fee to 696 *bps*. This is also reflected in the Sharpe ratios (*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*.

Figure 3 demonstrates an alternative way of measuring the impact of correlation timing to a dynamically optimizing investor by showing the cumulative wealth generated by the strategies. Over the 31-year sample with daily rebalancing, for the plug-in (Bayesian) allocation, an initial investment of \$1 grows to \$331 (\$310) for *MLR*, \$1,139 (\$1,016) for *CCC*, and \$3,064 (\$2,520) for *DCC*.²¹ The risk required to generate these returns appears to be quite reasonable. This is shown in Figure 4, which plots 3-year rolling estimates of the Sharpe ratios that for the most part range between 0 and 3. Clearly, therefore, it is worth using high-dimensional multivariate models for dynamic correlations as they generate significant economic value.

The exact value of the performance fees and the Sharpe ratios depends on the investor’s degree of relative risk aversion, λ . Since the tables focus on the case where $\lambda = 6$, the fees can be high. If we increase λ , we should expect the fees and Sharpe ratios to be lower; as investors become more risk averse they take less risk, and hence exploit less the correlation and volatility forecasts of a better

²¹Investing in the riskless asset (the yield of the US bond) generates a final wealth of \$8.6. The large difference between riskless investing and the benchmark static *MLR* model is due to the *MLR*’s optimal use of fluctuations in the US and foreign bonds.

model. Figure 5 illustrates an essential aspect of our findings by showing that the fees and Sharpe ratios are consistently higher with dynamic correlations as RRA increases from $\lambda = 6$ to $\lambda = 30$. Therefore, dynamic correlations are economically valuable even for large values of RRA in investor preferences.

Table 3 uncovers another important finding, which 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).²² This is not the case in evaluating correlation timing in the context of dynamic asset allocation strategies. 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. In our analysis, economic value is generated by making volatilities and correlations dynamic, irrespective of their exact specification. As we will see, this result holds for all three rebalancing frequencies we consider. 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 a model as any among the ones we consider. In our framework, increasing the sophistication of the econometric specification does not enhance the economic value of the simple *DCC* model.

Our analysis has so far focused on daily rebalancing for which the case for correlation timing is the strongest. We now turn to evaluating portfolio performance at the lower frequencies of weekly and monthly rebalancing. Since our forecasting exercise remains one-step ahead, we re-estimate all the models using weekly and monthly FX returns data. We begin with Table 4, which presents the results for weekly rebalancing and demonstrates that as we move from *MLR* to *CCC* we generate a performance fee of 158 *bps*, while the Sharpe ratio rises from 1.12 to 1.21. However, the *DCC* model now gives small incremental gains by slightly raising the *SR* to 1.28, while slightly lowering the fee to 122 *bps*.²³ This leads us to conclude that at the weekly frequency there is good economic value in dynamic volatilities but little value in dynamic correlations. In other words, it is economically valuable to have a dynamic covariance matrix; however, our strategies perform well even when correlations are constant and hence the dynamics of covariances are governed exclusively by the dynamics of volatilities. Furthermore, as shown in Table 5, at the monthly rebalancing there is very little or no economic value in both volatility and correlation timing.²⁴ Finally, the weakening economic evidence on correlation and volatility timing at these lower frequencies is in line with the

²²We 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.

²³This apparent inconsistency of a strategy raising the Sharpe ratio while lowering the fee is particular to the CRRA utility function and will be discussed in detail in the next section.

²⁴This is in fact consistent with Della Corte, Sarno and Tsiakas (2008), who find that the *CCC* – *GARCH* model generates little or no economic value for monthly FX strategies.

diminishing persistence of correlations and volatilities as we move from the daily to the monthly frequency.

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 correlation and volatility 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. We examine the effect of parameter uncertainty and Bayesian learning on the FX strategies with daily rebalancing in Panels B and C of Table 3. In Panel B we allow for parameter uncertainty only in the volatilities and correlations of daily returns, but not in their means. Panel C shows the results with full Bayesian learning.

Our main finding here is that parameter uncertainty in second moments has little or no effect on portfolio performance. Therefore, correlations and volatilities are estimated with high precision. This is a robust result that holds for all rebalancing frequencies, and for which there is a simple intuitive explanation: both volatilities and correlations are highly persistent and hence predictable, and so it makes sense that parameter uncertainty will not play a prominent role for a one-step ahead predictive horizon, irrespective of the frequency of the data.

It is also important to note that this is not the case for the first moments, which are notoriously difficult to estimate with high precision. Our results indicate that parameter uncertainty has a stronger impact on the mean than on correlations and volatilities. The prevailing view in the international finance literature is that exchange rates are not predictable, especially at short horizons, and follow a random walk (e.g., Mark, 1995). Our results here suggest that the constant in the random walk model for log-exchange rates (which is the same as the constant mean in the *MLR* model) is hard to pin-point with great accuracy. This is particularly true for daily FX returns. For example, the fee for the *DCC* model falls from 696 in the plug-in case to 655 in the full Bayesian learning case, whereas the Sharpe ratio falls from 1.62 to 1.58. As discussed in Barberis (2000), parameter uncertainty makes 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 enjoy lower fees and lower Sharpe ratios.²⁵ At any rate, despite the uncertainty over a large number of parameters, the economic gains from correlation timing remain strong.

²⁵In the framework of Barberis (2000), the additional risk of parameter uncertainty implies investing more in the riskless asset. In our case of multiple risky assets and allowing for negative weights, cutting the position on the risky assets does not necessarily imply a larger position on the riskless asset.

6.3.2 The Effect of Model Risk

We evaluate the impact of model risk on correlation timing by exploring whether portfolio performance improves when combining the forecasts stemming from the large set of models we estimate. We focus on the Bayesian Model Averaging (*BMA*) and Bayesian Model Winner (*BMW*) strategies, which provide for dynamic weights evaluated ex ante. The *BMA* and *BMW* strategies are applied on 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*).²⁶

The economic value of combined forecasts is reported in Table 6, which assesses the impact of model risk under the Bayesian predictive density. Therefore, we focus on the case of evaluating model risk across forecast combinations while at the same time accounting for parameter uncertainty and Bayesian learning. The results in Table 6 indicate that there is high economic value in *BMA* forecast combinations and even higher value in the *BMW* strategy for daily rebalancing. For instance, compared to the *DCC* – *GARCH* model, which under full Bayesian learning 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. However, accounting for model risk does not alter the previous finding that dynamic correlations generate little or no economic value at lower frequencies. In conclusion, therefore, accounting for model risk by forming combined forecasts delivers additional economic value and makes the case for correlation timing at the daily rebalancing frequency more robust.

6.3.3 Further Aspects of Asset Allocation

We have so far assessed the economic value of correlation timing in the FX market for an investor who is uncertain about the parameter estimates and the model specification. This section discusses two further aspects of the impact of dynamic return correlations on asset allocation: the effect of transaction costs and the effect of the choice of utility function.

We begin with an examination of the role of 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 benchmark strategy. A simple way of addressing this concern is by checking how much the performance fees decrease when we deduct a proportional

²⁶In this paper, we focus on forecast combinations of alternative specifications of *DCC* models with *GARCH* volatility. See Della Corte, Sarno and Tsiakas (2008) for the results on a forecast combination of *GARCH* and stochastic volatility models on monthly FX returns; and Liu and Maheu (2008) for Bayesian model averaging across univariate realized volatility models.

transaction cost of 1 *bps* or 2 *bps* from the portfolio return ex post. We carry out this calculation in each time period by first computing the optimal weights in the absence of any transaction costs. We then deduct from the portfolio return a transaction cost that is proportional to the dollar value traded in each risky asset. This approach ignores the fact that the dynamic portfolios are no longer optimal in the presence of transaction costs but maintains simplicity and tractability in our analysis. The results in Table 3 indicate that for daily rebalancing, the economic value of correlation timing is fairly robust to transaction costs. For a proportional daily transaction cost of 1 *bps* (2 *bps*) the performance fees decrease slightly from 386 *bps* to 349 (313) *bps* for *CCC – GARCH*, and from 696 *bps* to 637 (579) *bps* for *DCC – GARCH*.

In addition to computing performance fees net of transaction costs, we take a step further 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 basis points at the rebalancing frequency. For daily rebalancing, 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 can conclude that the economic value of correlation timing is robust to reasonable transaction costs for daily rebalancing.

In evaluating portfolio performance with weekly and monthly rebalancing, we observe that correlation timing may increase the Sharpe ratio, while decreasing the performance fee. This apparent inconsistency is specific to the use of CRRA utility. For this utility function, the optimal weights do depend on the mean-variance tradeoff, but the performance fees depend only on the portfolio returns generated by the weights. Since we optimally use the forecasts of each strategy to maximize utility (not the expected return of the portfolio), it is possible that a better strategy will generate portfolio returns with higher Sharpe ratio but with both lower expected returns and lower volatility. In this case, CRRA utility produces lower (or even negative) fees despite the higher Sharpe ratio.²⁷

This situation will not arise with quadratic utility, where the calculation of performance fees explicitly takes into account the mean-variance tradeoff. In fact, for weekly and monthly rebalancing, CRRA and quadratic utility will generate portfolios with almost identical expected returns, volatilities and Sharpe ratios; and yet, the performance fees will be very different across the two utility functions. This is not the case for daily rebalancing, when CRRA utility generates portfolios with much higher Sharpe ratios. Our analysis adopts CRRA utility because it is more general and does

²⁷Recall that the Sharpe ratio is an unconditional measure, which ignores higher moments.

not suffer from the well-known shortcomings of quadratic utility.²⁸ It is still, however, of considerable interest to evaluate the robustness of our findings to quadratic utility because quadratic utility is widely used and provides a high degree of analytical tractability in solving the dynamic allocation problem.²⁹

Table 7 reports the evidence on portfolio performance of selected models with quadratic utility. The results demonstrate that the CRRA and quadratic utility functions produce similar results, but the value of correlation timing tends to be a bit more pronounced with CRRA utility, particularly for daily rebalancing. The use of quadratic utility also confirms our main finding that correlation timing is most valuable with daily rebalancing. At weekly and monthly frequencies, the incremental value of dynamic correlations is small and dynamic volatilities are sufficient in capturing the dynamics of covariances.

6.4 Summary of Results

In the context of dynamic asset allocation strategies, our results provide strong evidence establishing the high economic value of timing short-horizon FX correlations. The value of correlation timing is distinct from that of volatility timing in that dynamic correlations add considerable economic gains to the case of dynamic volatilities with constant correlations. In short, we find that the economic value of correlation timing: (i) remains high when accounting for parameter uncertainty and Bayesian learning as investors are willing to accept only a slightly lower risk-return tradeoff; (ii) increases when accounting for model risk by optimally forming combined forecasts based on Bayesian Model Averaging; (iii) continues to be sizeable as RRA increases; (iv) remains largely unaffected when specifying asymmetric correlations, rich correlation dynamics, and alternative volatility specifications; (v) is robust to reasonable transaction costs in the FX market; (vi) is substantially high for daily rebalancing, but low to non-existent for weekly and monthly rebalancing; and finally (vii) it is similarly strong whether we use CRRA or quadratic utility.

7 Conclusion

The empirical finance literature has long determined that asset return correlations are dynamic. Therefore, accurate correlation forecasts are critical for investors' optimal asset allocation. This has motivated a long line of research dedicated to developing tractable multivariate volatility models. Despite the interest in dynamic correlations, however, the economic value of correlation timing has

²⁸If the utility function is quadratic, mean-variance analysis applies exactly and the allocation will not depend on higher moments. A well known weakness of quadratic utility is that it exhibits increasing RRA with respect to wealth.

²⁹In fact, the literature assessing the economic value of empirical models of asset returns predominantly uses quadratic utility (e.g., West, Edison and Cho, 1993; Fleming, Kirby and Ostdiek, 2001; Marquering and Verbeek, 2004; Han, 2006; Della Corte, Sarno and Thornton, 2008; and Della Corte, Sarno and Tsiakas, 2008).

not received as much attention as volatility timing. The justification for this partly lies on the substantial difficulty in handling multivariate models for dynamic correlations. At the same time, there appears to be a gap in the literature, as we know little about the portfolio implications of dynamic correlations in asset returns.

This paper bridges this gap by providing a comprehensive evaluation of how dynamic return correlations affect the optimal portfolio choice of risk-averse investors. We thus contribute to three related literatures in international finance, empirical asset pricing and financial econometrics in the following way. We design a new method for Bayesian estimation of the DCC model. Then, we use the dynamic correlation forecasts to formally assess the economic value of correlation timing. Our dynamic strategies invest in the FX market, and focus on the portfolio decision of an investor, who is uncertain about the parameter estimates and the model specification.

Our findings establish that there is high economic value in timing FX correlations, in addition to the substantial economic value in volatility timing. This result is robust to reasonable transaction costs, which in FX trading are generally low. It only holds, however, when rebalancing at the daily frequency since there is little or no economic value in timing FX correlations for weekly or monthly rebalancing. We also show 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 multivariate DCC model is a powerful instrument in the practice of currency management. In the end, correlation timing does matter to an international investor, and it pays to take dynamic FX correlations into consideration in asset allocation strategies when rebalancing at daily frequency.

As this study is the first to comprehensively assess the economic value of correlation timing, there is scope to potentially extend our analysis in several directions. For instance, given that our investigation is computationally demanding, it is perhaps worth exploring ways that make feasible the out-of-sample evaluation of dynamic correlation models in the context of asset allocation with Bayesian learning. A second direction is to extend the econometric models and the asset allocation to explicitly account for higher-order moments, such as skewness and kurtosis. Finally, it would be of interest to academics and practitioners alike to evaluate the effect of correlation timing for different asset classes, including stocks and bonds. We leave these extensions for future research.

Appendix A: The Set of Univariate Models for Volatility Timing

We estimate the multivariate correlation models under nine univariate volatility specifications:

1. GARCH: Bollerslev (1986):

$$\sigma_{i,t}^2 = \omega_i + \alpha_i u_{i,t-1}^2 + \beta_i \sigma_{i,t-1}^2.$$

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

$$\sigma_{i,t} = \omega_i + \alpha_i |u_{i,t-1}| + \beta_i \sigma_{i,t-1}.$$

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

$$\sigma_{i,t}^\tau = \omega_i + \alpha_i |u_{i,t-1}|^\tau + \beta_i \sigma_{i,t-1}^\tau.$$

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

$$\ln(\sigma_{i,t}^2) = \omega_i + \alpha_i \varepsilon_{i,t-1} + \kappa_i (|\varepsilon_{i,t-1}| - E[|\varepsilon_{i,t}|]) + \beta_i \ln(\sigma_{i,t-1}^2); \quad \varepsilon_{i,t-1} = \frac{u_{i,t-1}}{\sigma_{i,t-1}}; \quad E[|\varepsilon_{i,t}|] = \sqrt{\frac{2}{\pi}}.$$

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

$$\sigma_{i,t} = \omega_i + \alpha_i (|u_{i,t-1}| - \kappa_i u_{i,t-1}) + \beta_i \sigma_{i,t-1}.$$

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

$$\sigma_{i,t}^2 = \omega_i + \alpha_i (|u_{i,t-1}| - \kappa_i u_{i,t-1})^2 + \beta_i \sigma_{i,t-1}^2.$$

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

$$\sigma_{i,t}^\tau = \omega_i + \alpha_i (|u_{i,t-1}| - \kappa_i u_{i,t-1})^\tau + \beta_i \sigma_{i,t-1}^\tau.$$

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

$$\sigma_{i,t}^2 = \omega_i + \alpha_i (u_{i,t-1} + \kappa_i)^2 + \beta_i \sigma_{i,t-1}^2.$$

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

$$\sigma_{i,t}^2 = \omega_i + \alpha_i (u_{i,t-1} + \kappa_i \sigma_{i,t})^2 + \beta_i \sigma_{i,t-1}^2.$$

Appendix B: Bayesian MCMC Estimation

We perform Bayesian MCMC estimation 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. The MCMC algorithm for three representative multivariate models (*MLR*, *CCC*, *ADCC_{diag}*) is summarized below. All other specifications are special cases of one of these models.

The Multivariate Linear Regression (*MLR*) Algorithm

In the *MLR* model, we need to estimate $\theta = \{\mu, \Sigma^{-1}\}$, where μ is the vector of unconditional means, and Σ^{-1} is the constant precision matrix defined as the inverse of the covariance matrix. Given N assets, for μ we assume a Normal prior $N(\underline{m}, \underline{M})$, where $\underline{m} = 0_N$ and $\underline{M} = I_N$. For Σ^{-1} we specify a Wishart prior $W_N(\underline{s}, \underline{S})$, with scale matrix $\underline{S} = I_N$, degrees of freedom $\underline{s} = N + 4$, and mean $E[\Sigma^{-1}] = \underline{s}I_N$. The simple Gibbs algorithm is summarized below (for more details, see Koop, 2003):

1. Initialize Σ^{-1} .
2. Sample $\mu \mid y, \Sigma^{-1} \sim N_N(\bar{\mu}, \bar{M})$, where $\bar{M} = [\underline{M}^{-1} + \Sigma^{-1} \otimes (I' I)]^{-1}$,
and $\bar{\mu} = \bar{M} [\underline{M}^{-1} \underline{m} + (\Sigma^{-1} \otimes I') \text{vec}(y)]$.
3. Sample $\Sigma^{-1} \mid y, \mu \sim W_N(\bar{s}, \bar{S})$, where $\bar{s} = T + \underline{s}$ and $\bar{S} = [\underline{S}^{-1} + (y - \mu)'(y - \mu)]^{-1}$.
4. Go to step 2 and iterate 100,000 times beyond a burn-in of 20,000 iterations.

The Constant Conditional Correlation (*CCC*) Algorithm

In the *CCC* model with *GARCH* volatility (Bollerslev, 1990) we need to estimate $\theta = \{\mu, \theta_1, \bar{R}^{-1}\}$, where μ is the vector of unconditional means, $\theta_1 = \{\omega_i, \alpha_i, \beta_i\}$ are the conditional variance parameters for each of $i \leq N$ assets, and \bar{R}^{-1} is the inverse of the constant correlation matrix. The prior for μ is the same as for the *MLR* model above. For θ_1 we specify: $\omega \sim \text{LogN}(\underline{w}, \underline{W})$, with $\underline{w} = -4.8$ and $\underline{W} = 4$; $\alpha \sim \text{Beta}(\underline{a}, \underline{A})$, with $\underline{a} = 2$ and $\underline{A} = 40$; and $\beta \sim \text{Beta}(\underline{b}, \underline{B})$, with $\underline{b} = 40$ and $\underline{B} = 4$. These hyperparameters imply prior means: $E[\omega] = 0.06$, $E[\alpha] = 0.05$ and $E[\beta] = 0.89$, and hence $E[\sigma^2] = 1$. For \bar{R}^{-1} we assume a Wishart prior $W_N(\underline{s}, \underline{S})$, with scale matrix $\underline{S} = I_N$ and degrees of freedom $\underline{s} = N + 4$. The *CCC* algorithm is summarized below:

1. Initialize μ and transform the data into $u_t = (y_t - \mu)$.
2. For each univariate volatility process, sample the variance parameters θ_1 from their full conditional posterior density: $\theta_1 \mid y, \mu, \bar{R}^{-1}$. This posterior density is not available analytically. We compute the log-likelihood of the transformed data u_t as a function of θ_1 and then we optimize the log-posterior. We generate a proposal from a t -distribution $t(m, V, \xi)$, where m is the mode, V is the scaled inverse of the negative Hessian and ξ is a tuning parameter. The proposal is then accepted according to the independence chain Metropolis-Hasting algorithm (e.g., Chib and Greenberg, 1995).

3. Given the θ_1 draws from Step 2, we form the diagonal volatility matrix D_t and transform the data into $z_t = D_t^{-1}u_t = R_t^{1/2}\varepsilon_t$. Then, we sample \bar{R}^{-1} from a Wishart posterior distribution as in Step 3 of the *MLR* algorithm.
4. Sample the vector of unconditional means $\mu \mid y, \theta_1, \bar{R}^{-1}$ using a precision-weighted average of a set of normal priors and the normal likelihood conditional on θ_1 and \bar{R}^{-1} . Then, update the data $u_t = (y_t - \mu)$.
5. Go to step 2 and iterate 5,000 times beyond a burn-in of 1,000 iterations.

We design similar algorithms for alternative univariate volatility specifications.

The Asymmetric Dynamic Conditional Correlation (*ADCC_{diag}*) Algorithm

The asymmetric diagonal *ADCC_{diag}* model with *GARCH* volatility requires estimation of the parameter matrix $\theta = \{\mu, \theta_1, \theta_2\}$, where μ is the vector of unconditional means, $\theta_1 = \{\omega_i, \alpha_i, \beta_i\}$ are the *GARCH* parameters for each $i \leq N$ asset, and $\theta_2 = \{\Gamma, \Delta, \Pi\}$ are diagonal $N \times N$ matrices. In this model, we specify priors for μ as in the *MLR* model and priors for $\theta_1 = \{\omega_i, \alpha_i, \beta_i\}$ as in the *CCC* model above. For θ_2 , we specify multivariate Beta priors: $\Gamma \sim \text{Beta}(\underline{\gamma}, \underline{\Gamma})$, $\Delta \sim \text{Beta}(\underline{\delta}, \underline{\Delta})$, and $\Pi \sim \text{Beta}(\underline{\pi}, \underline{\Pi})$, where we set the scale matrices $\underline{\Gamma} = \underline{\Delta} = \underline{\Pi} = I_N$ and the degrees of freedom parameters $\underline{\gamma} = \underline{\delta} = \underline{\pi} = N + 1$. For all models, we set the hyperparameters to reasonable values, but the estimation algorithms are robust to prior specification and initial values. The *ADCC_{diag}* algorithm is summarized below:

1. Initialize μ and transform the data into $u_t = (y_t - \mu)$.
2. For each univariate volatility process, sample the conditional variance parameters θ_1 from their full conditional posterior density: $\theta_1 \mid u, \mu, \theta_2$ following Step 2 of the *CCC* algorithm described above. Use the θ_1 parameter estimates to form the diagonal volatility matrix D_t and transform the data into $z_t = D_t^{-1}u_t = R_t^{1/2}\varepsilon_t$.
3. Sample the conditional correlation parameters θ_2 from their full conditional posterior density: $\theta_2 \mid z, \mu, \theta_1$. This posterior density is not available analytically. Hence, we compute the log-likelihood of the transformed data z_t as a function of θ_2 and then we optimize the log-posterior. We generate a proposal from a *t*-distribution $t(m, V, \xi)$, where m is the mode, V is the scaled inverse of the negative Hessian and ξ is a tuning parameter. The proposal is then accepted according to the independence chain Metropolis-Hasting algorithm (e.g., Chib and Greenberg, 1995).
4. Sample the vector of unconditional means $\mu \mid y, \theta_1, \theta_2$ using a precision-weighted average of a set of normal priors and the normal likelihood conditional on θ_1 and θ_2 .
5. Go to step 2 and iterate 5,000 times beyond a burn-in of 1,000 iterations.

Appendix C: The Brandt *et al.* (2005) Solution to the Dynamic Asset Allocation Problem

In what follows we present the solution to the intertemporal asset allocation problem proposed by Brandt, Goyal, Santa-Clara and Stroud (2005). This is a simulation-based method for solving discrete-time portfolio choice problems, which allows us to directly use the correlation and volatility forecasts from the models in computing the dynamic weights. This method also makes it possible to extend our analysis to a Bayesian setting, where the conditional expectation is evaluated under each model's posterior predictive density.

Step 1: Expanding the Value Function

We can greatly simplify the problem by expanding the value function in a Taylor series around $W_t R^f$ and solve for the optimal weights. For the CRRA utility function, the solution to a second-order expansion of the value function is as follows:

$$x_t \approx \frac{1}{\lambda} \left\{ E_t \left[\psi_{t+1}(Z_{t+1}) \left((R_{t+1} - R^f) (R_{t+1} - R^f)' \right) \right] \right\}^{-1} \times E_t \left[\psi_{t+1}(Z_{t+1}) (R_{t+1} - R^f) \right].$$

This approximate solution for the optimal weights involves the conditional expectation of the second moment matrix (the variance-covariance matrix of demeaned excess returns) and the risk premium.

A fourth order expansion of the value function gives the following solution:

$$\begin{aligned} x_t \approx & \frac{1}{\lambda} \left\{ E_t \left[\psi_{t+1}(Z_{t+1}) \left((R_{t+1} - R^f) (R_{t+1} - R^f)' \right) \right] \right\}^{-1} \times E_t \left[\psi_{t+1}(Z_{t+1}) (R_{t+1} - R^f) \right] \\ & + \frac{1}{2} W_t^3 \partial^3 U(W_t R^f) E_t \left[\psi_{t+1}(Z_{t+1}) \left(x_t' (R_{t+1} - R^f) \right)^2 (R_{t+1} - R^f) \right] \\ & + \frac{1}{6} W_t^4 \partial^4 U(W_t R^f) E_t \left[\psi_{t+1}(Z_{t+1}) \left(x_t' (R_{t+1} - R^f) \right)^3 (R_{t+1} - R^f) \right], \end{aligned}$$

where $U(W) = \frac{W^{1-\lambda}}{1-\lambda}$, $\partial U(W) = W^{-\lambda} > 0$, $\partial^2 U(W) = -\lambda W^{-(\lambda+1)} < 0$, $\partial^3 U(W) = \lambda(\lambda+1) W^{-(\lambda+2)} > 0$, and $\partial^4 U(W) = -\lambda(\lambda+1)(\lambda+2) W^{-(\lambda+3)} < 0$. It is straightforward to solve this implicit expression for the optimal weights in practice. We start by computing the portfolio weights for the second-order expansion of the value function and take this to be an initial guess for the optimal weights. We then enter this guess into the right-hand side of the fourth-order expansion, obtain a new solution and repeat for a few iterations until the weights do not change anymore.

Step 2: Simulate Sample Paths

We use the multivariate models to simulate forward a large number of hypothetical sample paths for the vector of returns as follows:

$$y_t = \mu_t + \Sigma_t \varepsilon_t, \quad \varepsilon_t \sim N(0, I_N),$$

where at each time period t we sample $m = 1, \dots, M = 1,000$ artificial draws (simulated paths) from the distribution of ε_t using the one-step ahead forecasts for μ_t and Σ_t . We denote the realization of the returns and state variables at time s along the m path as $Y_s^m = (R_s^m, Z_s^m)$, where $Z_{t+1} = R_t$.

Step 3: Recursive Solution

We solve recursively backwards for the optimal portfolio weights at each date t for each simulated sample path m assuming that the portfolio weights from time $t + 1$ to $T - 1$ have already been computed and are denoted by \hat{x}_s , $s = t, \dots, T - 1$. We approximate terminal wealth as follows:

$$\widehat{W}_T = W_t R^f \prod_{s=t+1}^{T-1} \left(\hat{x}'_s (R_{s+1} - R^f) + R^f \right).$$

Consider now the problem of solving for the current portfolio weights given the current wealth W_t^m and the realization of the state variables Z_t^m . The weights for the second order expansion are:

$$\begin{aligned} x_t \approx & \left(-E_t \left\{ \partial^2 U \left(\widehat{W}_T \right) \prod_{s=t+1}^{T-1} \left(\hat{x}'_s (R_{s+1} - R^f) + R^f \right)^2 (R_{t+1} - R^f) (R_{t+1} - R^f)' \right\} W_t \right)^{-1} \\ & \times E_t \left[\partial U \left(\widehat{W}_T \right) \prod_{s=t+1}^{T-1} \left(\hat{x}'_s (R_{s+1} - R^f) + R^f \right) (R_{t+1} - R^f) \right]. \end{aligned}$$

Similarly, for the fourth-order expansion the weights are given by:

$$\begin{aligned} x_t \approx & \left(-E_t \left\{ \partial^2 U \left(\widehat{W}_T \right) \prod_{s=t+1}^{T-1} \left(\hat{x}'_s (R_{s+1} - R^f) + R^f \right)^2 (R_{t+1} - R^f) (R_{t+1} - R^f)' \right\} W_t \right)^{-1} \\ & \times E_t \left[\partial U \left(\widehat{W}_T \right) \prod_{s=t+1}^{T-1} \left(\hat{x}'_s (R_{s+1} - R^f) + R^f \right) (R_{t+1} - R^f) \right] \\ & + \frac{1}{2} E_t \left[\partial^3 U \left(\widehat{W}_T \right) \prod_{s=t+1}^{T-1} \left(\hat{x}'_s (R_{s+1} - R^f) + R^f \right)^3 (x'_t (R_{t+1} - R^f))^2 (R_{t+1} - R^f) \right] W_t^3 \\ & + \frac{1}{6} E_t \left[\partial^4 U \left(\widehat{W}_T \right) \prod_{s=t+1}^{T-1} \left(\hat{x}'_s (R_{s+1} - R^f) + R^f \right)^4 (x'_t (R_{t+1} - R^f))^3 (R_{t+1} - R^f) \right] W_t^4. \end{aligned}$$

Recursive Algorithm

1. Solve for the weights recursively backwards starting at $T - 1$.
 - (a) solve for 2nd order Taylor series expansion, and
 - (b) plug into 4th order Taylor series expansion.
2. Simulate the return sample paths using the conditional moments from the set of models we have estimated.
3. Approximate the terminal wealth for each sample path.
4. Go to Step 1 and iterate.

Table 1
Descriptive Statistics for Daily Exchange Rate Returns

Panel A: Percent Returns				
	<i>USD/GBP</i>	<i>USD/EUR</i>	<i>USD/CHF</i>	<i>USD/JPY</i>
<i>Mean</i>	−0.0004	0.0070	0.0095	0.0117
<i>Std Dev</i>	0.620	0.653	0.736	0.664
<i>Min</i>	−7.59	−5.87	−5.83	−4.15
<i>Max</i>	4.64	4.14	5.17	6.40
<i>Skewness</i>	−0.356	−0.005	0.083	0.609
<i>Kurtosis</i>	9.78	6.02	6.25	9.27
Panel B: Correlation Matrix				
	<i>USD/GBP</i>	<i>USD/EUR</i>	<i>USD/CHF</i>	<i>USD/JPY</i>
<i>USD/GBP</i>	1.00	0.557	0.495	0.336
<i>USD/EUR</i>	0.557	1.00	0.819	0.491
<i>USD/CHF</i>	0.495	0.819	1.00	0.543
<i>USD/JPY</i>	0.336	0.491	0.543	1.00

The table summarizes the descriptive statistics for the daily percent exchange rate returns. The returns are computed as the first difference of log-exchange rates. The data range from January 2, 1976 to December 29, 2006 for a sample size of 8,069 daily observations.

Table 2
Posterior Means for the $ADCC_{diag}$ -GARCH Model

Panel A: Daily FX Returns							
	μ	ω	α	β	γ	δ	π
<i>USD/GBP</i>	−0.0004 (0.00003)	0.0065 (0.00001)	0.9146 (0.00006)	0.0705 (0.00005)	0.9811 (0.00001)	0.1534 (0.00005)	0.1324 (0.00011)
<i>USD/EUR</i>	0.0071 (0.00003)	0.0042 (0.00001)	0.9278 (0.00005)	0.0651 (0.00004)	0.9780 (0.00001)	0.1979 (0.00003)	0.0479 (0.00013)
<i>USD/CHF</i>	0.0095 (0.00003)	0.0053 (0.00001)	0.9360 (0.00005)	0.0561 (0.00004)	0.9792 (0.00001)	0.1929 (0.00003)	0.0379 (0.00013)
<i>USD/JPY</i>	0.0117 (0.00003)	0.0049 (0.00001)	0.9256 (0.00005)	0.0667 (0.00004)	0.9872 (0.00001)	0.1439 (0.00003)	0.0351 (0.00014)
Panel B: Weekly FX Returns							
<i>USD/GBP</i>	−0.0026 (0.00014)	0.0622 (0.00015)	0.8945 (0.00015)	0.0745 (0.00011)	0.9231 (0.00007)	0.2135 (0.00014)	0.2676 (0.00022)
<i>USD/EUR</i>	0.0341 (0.00014)	0.0756 (0.00027)	0.8770 (0.00021)	0.0903 (0.00013)	0.9423 (0.00005)	0.2558 (0.00012)	0.0908 (0.00018)
<i>USD/CHF</i>	0.0462 (0.00016)	0.0937 (0.00031)	0.8919 (0.00019)	0.0743 (0.00012)	0.9519 (0.00004)	0.2548 (0.00009)	0.0826 (0.00019)
<i>USD/JPY</i>	0.0578 (0.00015)	0.0886 (0.00027)	0.8895 (0.00019)	0.0734 (0.00011)	0.9809 (0.00002)	0.1312 (0.00008)	0.1602 (0.00020)
Panel C: Monthly FX Returns							
<i>USD/GBP</i>	−0.0175 (0.00064)	0.1051 (0.00087)	0.9312 (0.00020)	0.0583 (0.00017)	0.8594 (0.00041)	0.1911 (0.00036)	0.1600 (0.00074)
<i>USD/EUR</i>	0.1383 (0.00067)	0.2889 (0.00244)	0.9270 (0.00027)	0.0470 (0.00016)	0.8370 (0.00030)	0.2946 (0.00034)	0.2760 (0.00076)
<i>USD/CHF</i>	0.1907 (0.00072)	0.3592 (0.00276)	0.9026 (0.00031)	0.0714 (0.00021)	0.8779 (0.00024)	0.2950 (0.00030)	0.2286 (0.00067)
<i>USD/JPY</i>	0.2397 (0.00069)	0.4780 (0.00233)	0.9019 (0.00028)	0.0592 (0.00020)	0.9631 (0.00083)	0.0652 (0.00025)	0.3152 (0.00064)

The table presents the Bayesian MCMC estimates of the posterior means for the parameters of the diagonal asymmetric $ADCC_{diag}$ -GARCH model. This is the most general multivariate correlation specification we consider in our analysis. The model is applied on daily, weekly and monthly percent exchange rate returns for the sample period of January 2, 1976 to December 29, 2006. The numbers in parenthesis indicate the Numerical Standard Error (NSE).

Table 3
Portfolio Performance with Daily Rebalancing

Panel A: Plug-In Method							
Model	μ_p	σ_p	SR	$\Phi_{(\tau=0)}$	$\Phi_{(\tau=1)}$	$\Phi_{(\tau=2)}$	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	18.8	11.2	1.08				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	22.7	11.8	1.36	386	349	313	10.5
<i>CCC – AVGARCH</i>	22.1	12.0	1.28	327	281	234	7.0
<i>CCC – NARCH</i>	22.4	11.9	1.32	356	314	271	8.4
<i>CCC – EGARCH</i>	22.9	12.1	1.34	404	357	309	8.5
<i>CCC – ZARCH</i>	22.6	12.1	1.32	377	331	284	8.1
<i>CCC – GJR – GARCH</i>	22.8	11.8	1.37	399	363	326	10.9
<i>CCC – APGARCH</i>	22.7	11.9	1.34	383	341	299	9.1
<i>CCC – AGARCH</i>	23.0	11.8	1.37	412	375	338	11.2
<i>CCC – NAGARCH</i>	23.0	11.8	1.38	416	380	343	11.3
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	25.8	11.7	1.62	696	637	579	11.9
<i>DCC_{diag} – GARCH</i>	25.7	11.7	1.62	691	629	567	11.1
<i>ADCC – GARCH</i>	25.7	11.7	1.62	691	633	574	11.8
<i>ADCC_{diag} – GARCH</i>	25.8	11.7	1.62	694	631	569	11.1

(continued)

Table 3 (*continued*)
Portfolio Performance with Daily Rebalancing

Panel B: Bayesian Learning in Volatilities and Correlations							
Model	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	18.8	11.2	1.08				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	22.7	11.8	1.36	386	349	313	10.5
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	25.7	11.7	1.62	692	633	575	11.8
Panel C: Full Bayesian Learning in the Mean, Volatilities and Correlations							
Model	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	18.5	10.8	1.08				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	22.3	11.4	1.36	371	334	296	9.9
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	25.2	11.6	1.58	655	596	536	11.0

The table shows the portfolio performance of dynamic models against the static benchmark for daily rebalancing. *MLR* is the multivariate linear regression model, *CCC* is the constant conditional correlation model, and *DCC* is the dynamic conditional correlation model. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p , σ_p and SR respectively. The dynamic strategies use the maximum expected utility rule to 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. Panel A uses the plug-in estimates of the mean, volatility and correlation forecasts, Panel B evaluates expected utility using the volatility and correlation forecasts from the predictive density, and Panel C accounts for full Bayesian learning. The performance fee (Φ) denotes the amount 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 models. The performance fees are expressed in annual basis points and are reported for three levels of proportional transaction costs $\tau = \{0, 1, 2\}$ bps. 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 expressed in daily basis points.

Table 4
Portfolio Performance with Weekly Rebalancing

Panel A: Plug-In Method							
Model	μ_p	σ_p	SR	$\Phi_{(\tau=0)}$	$\Phi_{(\tau=1)}$	$\Phi_{(\tau=2)}$	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	20.6	12.3	1.12				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	22.2	12.8	1.21	158	150	143	19.9
<i>CCC – AVGARCH</i>	22.0	13.3	1.15	127	117	108	13.8
<i>CCC – NARCH</i>	22.4	13.0	1.20	169	161	153	20.0
<i>CCC – EGARCH</i>	22.3	13.5	1.15	156	144	133	13.8
<i>CCC – ZARCH</i>	22.2	13.5	1.14	144	134	124	14.1
<i>CCC – GJR – GARCH</i>	22.3	12.9	1.20	161	153	145	20.4
<i>CCC – APGARCH</i>	22.8	13.3	1.20	205	195	185	20.8
<i>CCC – AGARCH</i>	22.4	13.0	1.21	176	167	159	20.9
<i>CCC – NAGARCH</i>	22.4	13.0	1.21	174	166	157	20.8
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	21.7	11.7	1.28	122	113	105	14.4
<i>DCC_{diag} – GARCH</i>	22.0	12.1	1.26	144	132	119	11.7
<i>ADCC – GARCH</i>	21.7	11.7	1.28	121	113	104	14.3
<i>ADCC_{diag} – GARCH</i>	22.7	12.5	1.28	209	194	180	14.4

(continued)

Table 4 (*continued*)
Portfolio Performance with Weekly Rebalancing

Panel B: Bayesian Learning in Volatilities and Correlations							
Model	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	20.6	12.3	1.12				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	22.2	12.8	1.21	155	148	141	19.9
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	21.7	11.6	1.28	119	110	102	14.2
Panel C: Full Bayesian Learning in the Mean, Volatilities and Correlations							
Model	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	20.5	12.1	1.12				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	22.0	12.5	1.21	153	146	138	19.9
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	21.5	11.4	1.27	116	107	100	14.5

The table shows the portfolio performance of dynamic models against the static benchmark for weekly rebalancing. *MLR* is the multivariate linear regression model, *CCC* is the constant conditional correlation model, and *DCC* is the dynamic conditional correlation model. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p , σ_p and SR respectively. The dynamic strategies use the maximum expected utility rule to build a portfolio by investing in the 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. Panel A uses the plug-in estimates of the mean, volatility and correlation forecasts, Panel B evaluates expected utility using the volatility and correlation forecasts from the predictive density, and Panel C accounts for full Bayesian learning. The performance fee (Φ) denotes the amount 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 models. The performance fees are expressed in annual basis points and are reported for three levels of proportional transaction costs $\tau = \{0, 1, 2\}$ bps. The break-even transaction cost τ^{BE} is defined as the minimum weekly proportional cost that cancels out the utility advantage of a given strategy and is expressed in weekly basis points.

Table 5
Portfolio Performance with Monthly Rebalancing

Panel A: Plug-In Method							
Model	μ_p	σ_p	SR	$\Phi_{(\tau=0)}$	$\Phi_{(\tau=1)}$	$\Phi_{(\tau=2)}$	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	19.4	14.0	0.90				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	17.5	11.6	0.92	–160	–160	–160	–
<i>CCC – AVGARCH</i>	17.8	12.1	0.91	–137	–137	–137	–
<i>CCC – NARCH</i>	17.6	11.8	0.92	–145	–145	–145	–
<i>CCC – EGARCH</i>	18.2	12.5	0.91	–97	–98	–98	–
<i>CCC – ZARCH</i>	17.7	12.0	0.91	–142	–142	–142	–
<i>CCC – GJR – GARCH</i>	17.4	11.6	0.91	–169	–169	–170	–
<i>CCC – APGARCH</i>	17.6	11.7	0.92	–145	–145	–145	–
<i>CCC – AGARCH</i>	17.2	11.3	0.92	–182	–182	–183	–
<i>CCC – NAGARCH</i>	17.5	11.5	0.93	–154	–154	–155	–
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	17.3	10.7	0.98	–172	–172	–172	–
<i>DCC_{diag} – GARCH</i>	17.4	11.1	0.96	–163	–164	–165	–
<i>ADCC – GARCH</i>	17.4	10.9	0.98	–162	–162	–163	–
<i>ADCC_{diag} – GARCH</i>	17.6	11.5	0.94	–149	–151	–153	–

(continued)

Table 5 (*continued*)
Portfolio Performance with Monthly Rebalancing

Panel B: Bayesian Learning in Volatilities and Correlations							
Model	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	19.4	14.0	0.90				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	17.4	11.5	0.92	–164	–164	–165	–
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	17.2	10.7	0.98	–176	–176	–177	–
Panel C: Full Bayesian Learning in the Mean, Volatilities and Correlations							
Model	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>Static Benchmark</i>							
<i>MLR</i>	19.2	13.8	0.90				
<i>Volatility Timing</i>							
<i>CCC – GARCH</i>	17.4	11.4	0.92	–167	–167	–167	–
<i>Correlation and Volatility Timing</i>							
<i>DCC – GARCH</i>	17.0	10.5	0.97	–181	–182	–182	–

The table shows the portfolio performance of dynamic models against the static benchmark for monthly rebalancing. *MLR* is the multivariate linear regression model, *CCC* is the constant conditional correlation model, and *DCC* is the dynamic conditional correlation model. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p , σ_p and SR respectively. The dynamic strategies use the maximum expected utility rule to 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. Panel A uses the plug-in estimates of the mean, volatility and correlation forecasts, Panel B evaluates expected utility using the volatility and correlation forecasts from the predictive density, and Panel C accounts for full Bayesian learning. The performance fee (Φ) denotes the amount 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 models. The performance fees are expressed in annual basis points and are reported for three levels of proportional transaction costs $\tau = \{0, 1, 2\}$ bps. The break-even transaction cost τ^{BE} is defined as the minimum monthly proportional cost that cancels out the utility advantage of a given strategy and is expressed in monthly basis points.

Table 6
The Economic Value of Combined Forecasts with Bayesian Learning

Panel A: Daily Rebalancing							
Model Combination	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>BMA – VOL</i>	24.2	10.6	1.65	568	496	424	7.9
<i>BMA – CORR</i>	25.1	11.6	1.58	653	592	530	10.6
<i>BMA – FULL</i>	24.0	10.0	1.73	554	477	401	7.2
<i>BMW – VOL</i>	24.8	10.4	1.73	631	447	262	3.4
<i>BMW – CORR</i>	25.3	11.5	1.60	664	583	503	8.2
<i>BMW – FULL</i>	23.8	9.4	1.81	539	298	57	2.2
Panel B: Weekly Rebalancing							
<i>BMA – VOL</i>	21.5	11.2	1.31	112	104	96	13.4
<i>BMA – CORR</i>	21.8	11.5	1.31	140	131	122	15.8
<i>BMA – FULL</i>	21.5	11.1	1.33	117	109	100	13.5
<i>BMW – VOL</i>	21.4	11.2	1.30	100	75	50	4.0
<i>BMW – CORR</i>	22.0	11.2	1.36	162	149	136	12.4
<i>BMW – FULL</i>	22.0	10.6	1.44	174	138	101	4.8
Panel C: Monthly Rebalancing							
<i>BMA – VOL</i>	17.0	10.6	0.97	–187	–187	–187	–
<i>BMA – CORR</i>	17.1	10.5	0.98	–172	–172	–172	–
<i>BMA – FULL</i>	17.1	10.5	0.98	–180	–180	–180	–
<i>BMW – VOL</i>	17.4	10.5	1.01	–150	–153	–155	–
<i>BMW – CORR</i>	17.4	10.3	1.03	–148	–148	–148	–
<i>BMW – FULL</i>	18.2	10.6	1.08	–63	–67	–72	–

The table assesses the impact of model risk on correlation timing by presenting the portfolio performance of combined forecasts. Expected utility is evaluated under the predictive density thus accounting for Bayesian learning. *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, volatility and Sharpe ratio of each portfolio are denoted by μ_p , σ_p and SR respectively. The performance fee (Φ) denotes the amount 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. The performance fees are expressed in annual basis points and are reported for three levels of proportional transaction costs $\tau = \{0, 1, 2\}$ bps. 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 expressed in basis points at the rebalancing frequency.

Table 7
Portfolio Performance of Selected Models with Quadratic Utility

Panel A: Daily Rebalancing							
Model	μ_p	σ_p	SR	Φ ($\tau=0$)	Φ ($\tau=1$)	Φ ($\tau=2$)	τ^{BE}
<i>Plug-in case</i>							
<i>MLR</i>	19.9	13.3	0.99				
<i>CCC – GARCH</i>	24.0	14.5	1.19	305	258	211	6.5
<i>DCC – GARCH</i>	27.5	15.7	1.32	547	469	392	7.0
<i>Full Bayesian Learning</i>							
<i>MLR</i>	19.7	13.1	0.99				
<i>CCC – GARCH</i>	23.2	14.6	1.13	230	184	139	5.0
<i>DCC – GARCH</i>	26.7	15.8	1.26	467	390	313	6.1
Panel B: Weekly Rebalancing							
<i>Plug-in case</i>							
<i>MLR</i>	20.6	12.3	1.12				
<i>CCC – GARCH</i>	22.2	12.8	1.21	129	121	114	17.3
<i>DCC – GARCH</i>	21.7	11.6	1.28	163	154	146	19.1
<i>Full Bayesian Learning</i>							
<i>MLR</i>	20.4	12.1	1.12				
<i>CCC – GARCH</i>	22.0	12.5	1.21	127	120	112	16.9
<i>DCC – GARCH</i>	21.5	11.4	1.27	160	152	148	19.2
Panel C: Monthly Rebalancing							
<i>Plug-in case</i>							
<i>MLR</i>	19.2	13.9	0.90				
<i>CCC – GARCH</i>	17.4	11.5	0.92	10	10	10	9.2
<i>DCC – GARCH</i>	17.2	10.6	0.98	51	51	51	11.6
<i>Full Bayesian Learning</i>							
<i>MLR</i>	19.1	13.7	0.90				
<i>CCC – GARCH</i>	17.5	11.8	0.91	4	4	4	–
<i>DCC – GARCH</i>	18.3	11.0	0.96	35	35	35	9.5

The table evaluates the portfolio performance of selected models under quadratic utility and compares the plug-in case with full Bayesian learning. *MLR* is the multivariate linear regression model, *CCC* is the constant conditional correlation model, and *DCC* is the scalar symmetric dynamic conditional correlation model. The annualized percent mean, volatility and Sharpe ratio of each portfolio are denoted by μ_p , σ_p and SR respectively. The dynamic portfolio strategies invest in five bonds from the US, UK, Germany, Switzerland and Japan and use the four exchange rate forecasts to convert the portfolio return in US dollars. The performance fee (Φ) denotes the amount an investor with quadratic utility and a degree of relative risk aversion of 6 will pay to switch from *MLR* to one of the dynamic models. The performance fee is expressed in annual basis points and is reported for proportional transaction costs $\tau = \{0, 1, 2\}$ bps. The break-even transaction cost τ^{BE} is the minimum proportional cost that cancels out the utility advantage of a strategy and is expressed in basis points at the rebalancing frequency.

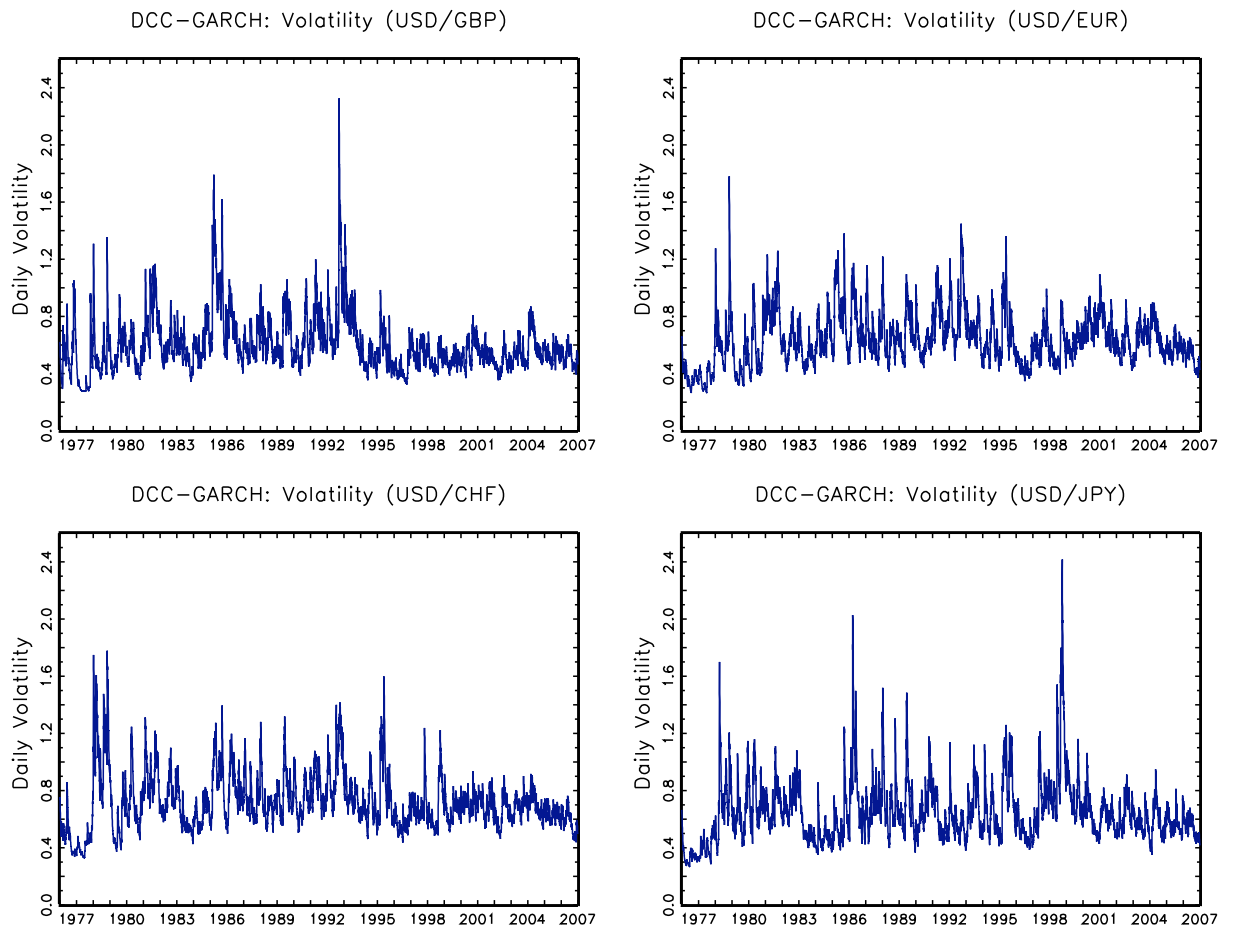


Figure 1: The daily volatility forecasts for the four exchange rate return series using the simple (scalar symmetric) DCC-GARCH model.

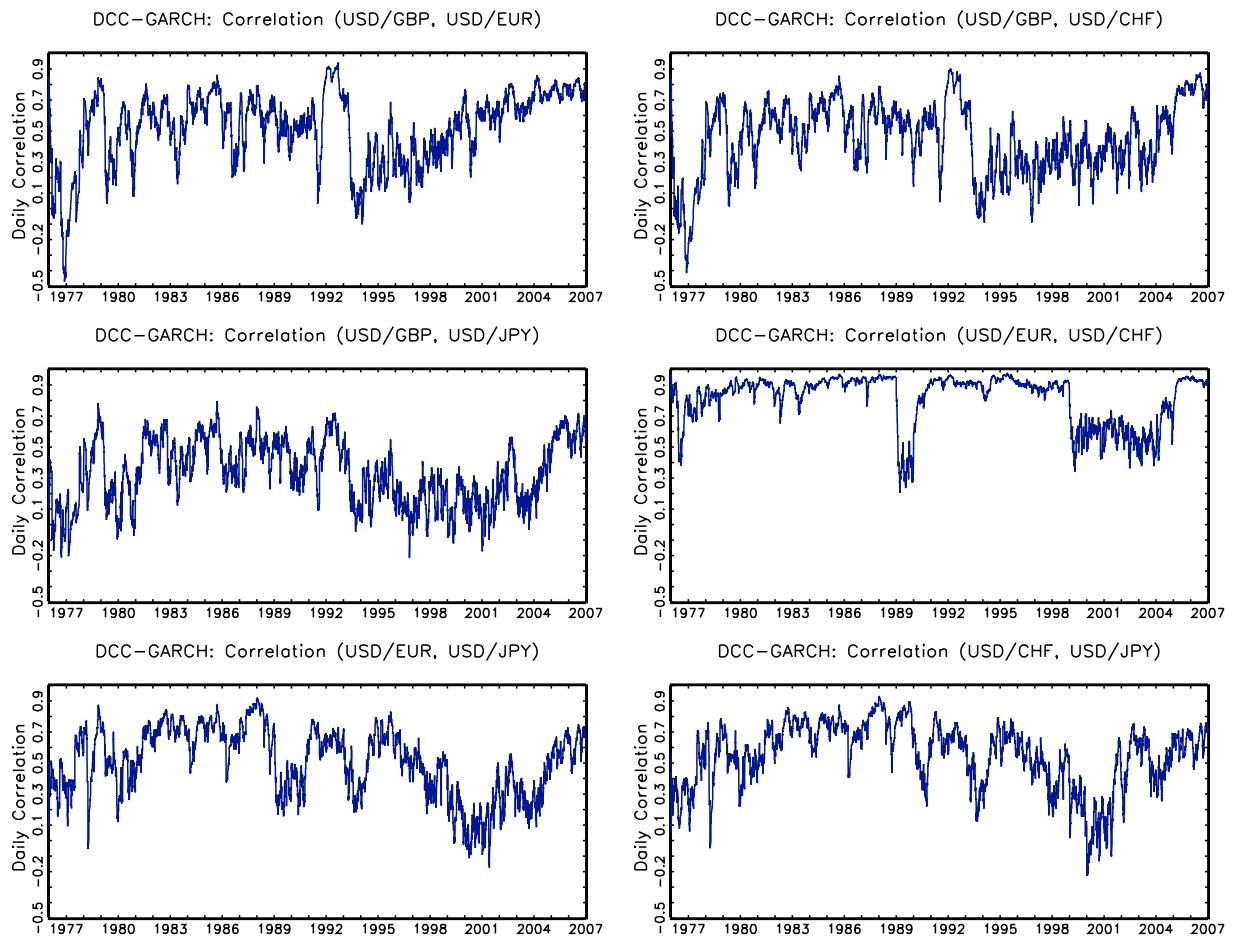


Figure 2: The daily correlation forecasts among the four exchange rate return series using the simple (scalar symmetric) DCC-GARCH model.

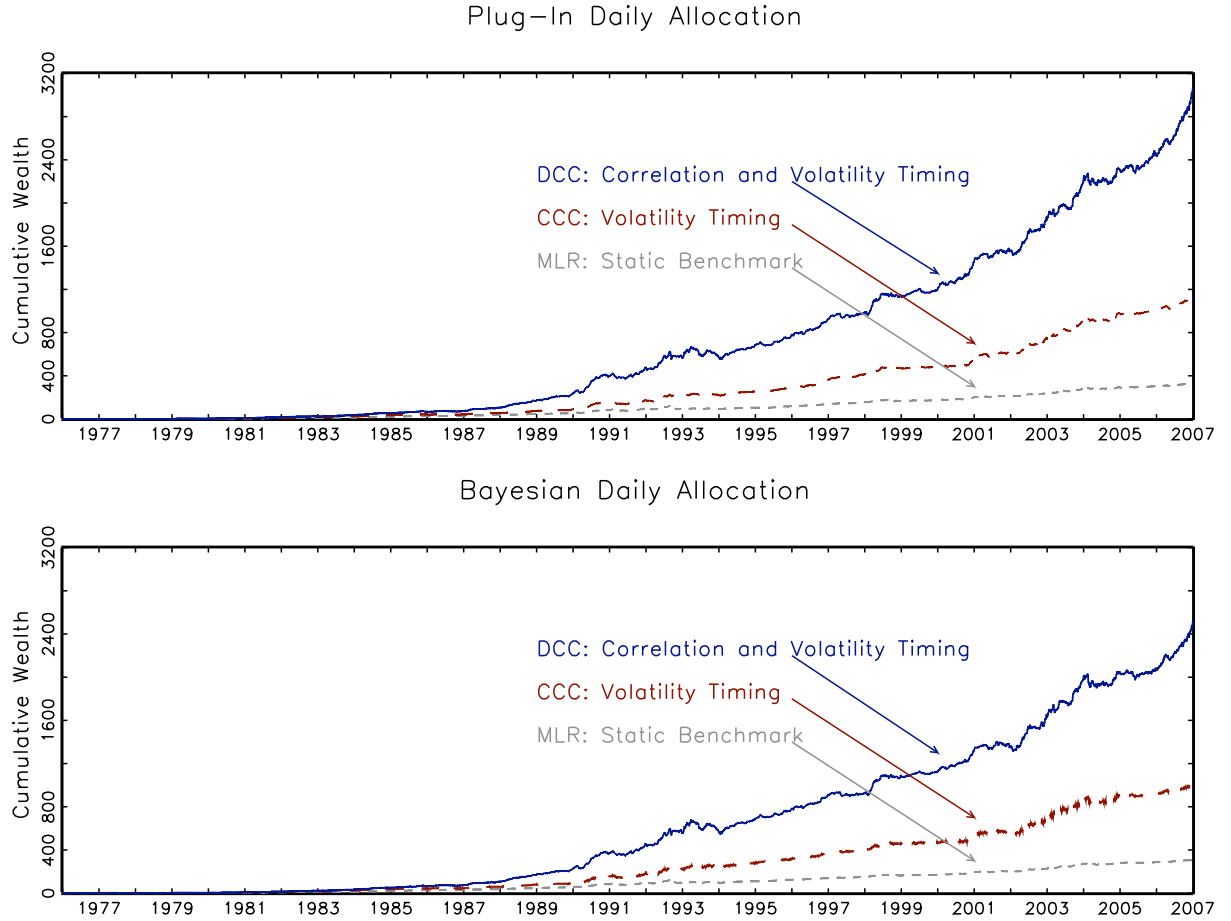


Figure 3: The cumulative wealth of dynamic strategies with daily rebalancing. Initial wealth is \$1 growing at the portfolio return. MLR is the benchmark model with static volatilities and correlations; CCC has dynamic GARCH volatilities, but constant correlations; and DCC has dynamic (scalar symmetric) correlations and GARCH volatilities. The top panel is for plug-in allocation and the bottom panel for Bayesian allocation. All strategies are evaluated for a coefficient of relative risk aversion equal to 6. The figure displays the case of zero transaction costs.

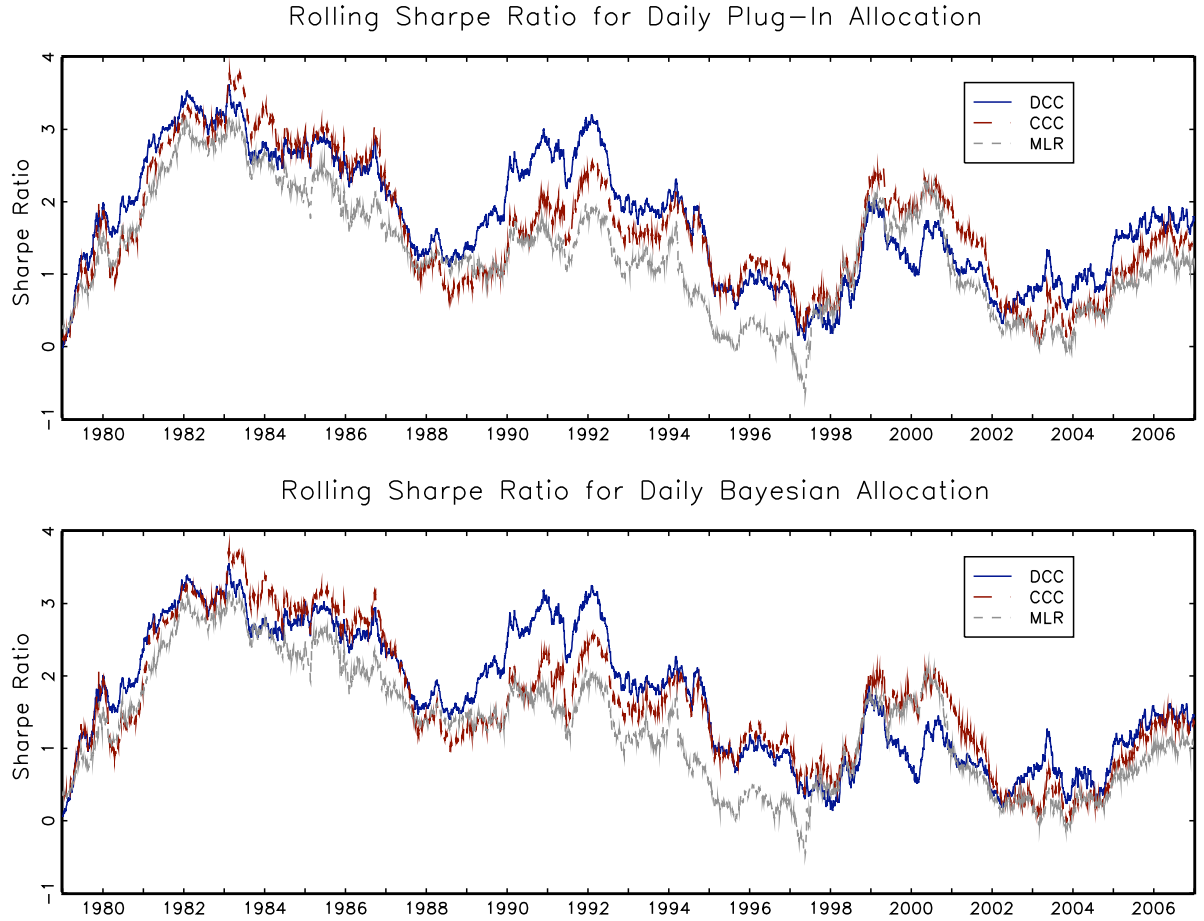


Figure 4: The Sharpe ratio of dynamic strategies with daily rebalancing. The Sharpe ratios are calculated using a rolling 3-year window. MLR is the benchmark model with static volatilities and correlations; CCC has dynamic GARCH volatilities, but constant correlations; and DCC has dynamic (scalar symmetric) correlations and GARCH volatilities. The top panel is for plug-in allocation and the bottom panel for Bayesian allocation. All strategies are evaluated for a degree of relative risk aversion equal to 6. The figure displays the case of zero transaction costs.

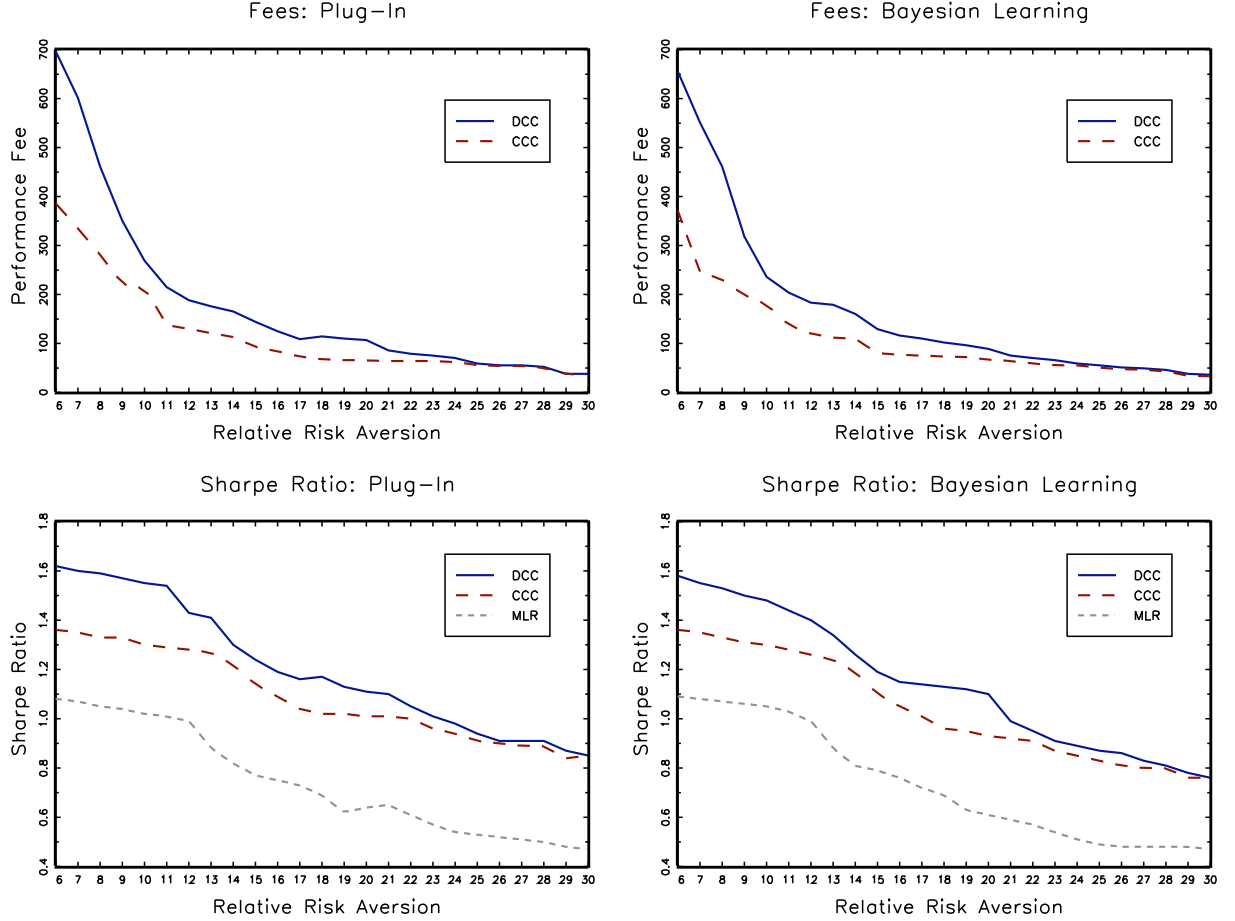


Figure 5: The performance of dynamic strategies with respect to the degree of relative risk aversion for daily rebalancing. MLR is the benchmark model with static volatilities and correlations; CCC has dynamic GARCH volatilities, but constant correlations; and DCC has dynamic (scalar symmetric) correlations and GARCH volatilities. The figure shows the performance fees for plug-in asset allocation (upper left panel), the performance fees for Bayesian allocation (upper right panel), the Sharpe ratios for plug-in allocation (lower left panel), and the Sharpe ratios for Bayesian allocation (lower right panel). The strategies are evaluated for a degree of relative risk aversion ranging from 6 to 30. The figure displays the case of zero transaction costs.

References

- Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2003). "Modeling and Forecasting Realized Volatility," *Econometrica* **71**, 579–625.
- 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.
- Ball, C.A, and W.N. Torous (2000). "Stochastic Correlation across International Stock Markets," *Journal of Empirical Finance* **7**, 373–388.
- Bank of International Settlements (2007). "Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity in April 2007," Monetary and Economic Department.
- 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.
- 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 (forthcoming).
- Brandt, M.W., and F.X. Diebold (2006). "A No-Arbitrage Approach to Range-Based Estimation of Return Covariances and Correlations," *Journal of Business* **79**, 61–73.
- 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, (2006). "Correlation Risk and Optimal Portfolio Choice," Unpublished working paper: Tanaka Business School.
- Campbell, J.Y., M. Lettau, B.G. Malkiel, and Y. Xu (2001). "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance* **56**, 1–43.

- 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. (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.
- Cochrane, J.H. (2008). "The Dog That Did Not Bark: A Defense of Return Predictability," *Review of Financial Studies* **21**, 1533–1575.
- Cremers, M. (2002). "Stock Return Predictability: A Bayesian Model Selection Perspective," *Review of Financial Studies* **15**, 1223–1249.
- Das, S.R., and R. Uppal (2004). "Systemic Risk and International Portfolio Choice," *Journal of Finance* **59**, 2809–2834.
- 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 (2008). "An Economic Evaluation of Empirical Exchange Rate Models." *Review of Financial Studies* (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 (2008). "The Price of Correlation Risk: Evidence from Equity Options," *Journal of Finance* (forthcoming).
- 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., N. Mark, and K.D. West (2007). "Exchange Rate Models Are Not As Bad As You Think," in D. Acemoglu, K. Rogoff and M. Woodford (eds.), *NBER Macroeconomics Annual* **22**, 381–441.
- 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., and K. Sheppard (2001). “Theoretical and Empirical Properties of Dynamic Conditional Correlation Multivariate GARCH,” Unpublished working paper, University of California at San Diego.
- Erb, C., C. Harvey, and T. Viskanta (1994). “Forecasting International Equity Correlations,” *Financial Analysts Journal* **50**, 32–45.
- Fiorentini, G., E. Sentana, and N. Shephard (2004). “Likelihood-Based Estimation of Latent Generalized ARCH Structures,” *Econometrica* **72**, 1481–1517.
- Fleming, J., C. Kirby, and B. Ostdiek (2001). “The Economic Value of Volatility Timing,” *Journal of Finance* **56**, 329–352.
- Geweke, J. (1989a). “Exact Predictive Densities for Linear Models with ARCH Disturbances,” *Journal of Econometrics* **40**, 63–86.
- Geweke, J. (1989b). “Bayesian Inference in Econometric Models Using Monte Carlo Integration,” *Econometrica* **57**, 1317–1339.
- 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 (2008). “The Wishart Autoregressive Process of Multivariate Stochastic Volatility,” *Journal of Econometrics* (forthcoming)
- 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.
- Hoeting, J.A., D. Madigan, A.E. Raftery, and C.T. Volinsky (1999). “Bayesian Model Averaging: A Tutorial,” *Statistical Science* **14**, 382–417.
- Inoue, A., and L. Kilian (2004). “In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?” *Econometric Reviews* **23**, 371–402.
- Inoue, A., and L. Kilian (2006). “On the Selection of Forecasting Models,” *Journal of Econometrics* **130**, 273–306.
- Jondeau, E., and M. Rockinger (2008). “The Economic Value of Distributional Timing,” Unpublished Working paper, University of Lausanne.

- 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.
- Koop, G. (2003). *Bayesian Econometrics*, New York: Wiley.
- Krishnan, C.N.V., R. Petkova, and P. Ritchken (2008). "Correlation Risk," Unpublished working paper, Case Western Reserve University.
- Leamer, E.E. (1978). *Specification Searches*, New York: Wiley.
- 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.
- Liu, C., and J.M. Maheu (2008). "Forecasting Realized Volatility: A Bayesian Model Averaging Approach," *Journal of Applied Econometrics* (forthcoming).
- 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.
- Mark, N.C. (1995). "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review* **85**, 201–218.
- 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.
- Nelson, D.B. (1991). "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica* **59**, 347–370.
- Philipov, A., and M.E. Glickman (2006). "Multivariate Stochastic Volatility via Wishart Processes," *Journal of Business and Economic Statistics* **24**, 313–328.
- Skoulakis, G. (2007). "Dynamic Portfolio Choice with Bayesian Learning," Unpublished working paper, University of Maryland.
- 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.
- 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.
- 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.