

Abstract

In this paper, we evaluate the economic value that arise from incorporating conditional volatility when forecasting the covariance matrix of returns for both short and long horizons in the Vietnamese stock market, using the volatility timing framework of Fleming et al. (2001). We report three main findings. First, investors are willing to pay to switch from the static to a dynamic volatility timing strategy. Second, there is negligible difference in forecast performance among short and memory volatility models. However, the more parsimonious EWMA family models tend to produce better forecasts of the covariance matrix than those produced by the GARCH family volatility models at all investment horizons. Third, when transaction costs are taken into account, the gains from daily rebalanced dynamic portfolios deteriorate. However, it is still worth implementing the dynamic strategies at lower rebalancing frequencies. Our results are robust to estimation error in expected returns, the choice of risk aversion coefficient and estimation windows.

Keywords: Conditional variance-covariance matrix; Volatility timing; Asset allocation; Economic value; Vietnamese stock markets.

JEL code: G11, G17, C58

Date of receipt: 29th Nov 2016; *Date of revision:* 27th December 2016; *Date of approval:* 30th Dec 2016

1 Introduction

Extensive research suggests that multivariate conditional volatility models produce better forecasts of the covariance matrix than those produced by the unconditional covariance matrix estimator (see, for example, Engle and Colacito, 2006). Exploiting the predictability of volatility and covariance has become a key driver in many applied areas of finance, including asset allocation, asset pricing and risk management. Fleming et al. (2001) are among the first to study the economic value of predicting and timing volatility for risk averse investors in an asset allocation setting. Expected returns are treated as constant and investors periodically update their

¹ PhD, Foreign Trade University, Hochiminh City Campus, Viet nam. Corresponding author, E-mail: anhth.nguyen@ftu.edu.vn

portfolios based on forecasts of the conditional covariance matrix. They show that investors are better off in terms of utility when switching from a static strategy based on the unconditional volatility estimator to dynamic volatility timing strategies based on conditional volatility models. Recent studies incorporate more properties of volatility dynamics in application to investment decisions. Thorpe and Milunovich (2007) allow for asymmetries in modelling volatility and correlation, and show that investors are willing to pay to switch from symmetric to asymmetric forecasts. Similarly, Hyde et al. (2010) demonstrate the benefits of accounting for volatility jumps in asset allocation strategies. In the dynamic economic value studies, the conditional covariance matrix is typically estimated applying popular conditional volatility models such as the multivariate Exponentially Weighted Moving Average (EWMA) or multivariate memory Generalised Autoregressive Conditional Heteroscedasticity (GARCH) models, where shocks to volatility and covariance dissipate rapidly due to their exponential weighting. Consequently, most of the studies on the economic value of the short memory conditional covariance matrix focus on short horizon day traders. While this approach may make the most use of the forecast power of the short memory conditional volatility models, it may not nevertheless correspond to the needs of most practical investors, who often rebalance their portfolios at lower frequencies.

A mounting body of empirical evidence now suggests that a shock to return volatility is more persistent than what is implied in the exponential decay of the short memory GARCH or EWMA models. This “long memory” feature is important not only for the measurement of current volatility, but also for forecasts of future volatility, especially over longer horizons. This has prompted the development of volatility models that incorporate long memory in volatility dynamics (for example, the Fractionally Integrated FIGARCH model of Baillie, Bollerslev and Mikkelsen (1996), the Hyperbolic HYGARCH model of Davidson (2004), the Component CGARCH model of Engle and Lee (1999) or the Long memory EWMA model of Zumbach (2006)). Long memory volatility models have been suggested to provide better estimates and forecasts of volatility than those generated by short memory volatility in both univariate and multivariate context (Teyssiere, 1998, Niguez and Rubia, 2006, Harris and Nguyen, 2013).

Accurate forecast of return volatility and covariance is of crucial importance in emerging markets where stock market volatility is much higher than that in developed markets. As one of the most rapidly developing emerging markets in the world, the Vietnamese stock market has attracted the interests of both investors and researchers. However, with only 15 years of development, the Vietnamese stock market is characterized by extreme stock return movements. Initiating in 2000, the Vietnamese major stock index, the VN index, reached its peak in March 2007 at around 1170 before losing 80% of its value to its lowest level of 280 in December 2008.

Understanding the sources and the characteristics of volatility in Vietnam is thus important for policymakers as well as investors. However, volatility modelling and forecasting in Vietnam has not attracted the deserved attention possibly because the stock market is largely underdeveloped. Indeed, so far the number of studies on volatility in Vietnam is limited and has generally restricted themselves to the analysis of volatility in a univariate setting (see, e.g., Tran Manh Tuyen, 2011, Vo Xuan Vinh and Nguyen Thi Kim Ngan, 2011). There is (to our best knowledge) still a lack of research that examines the benefits of allowing for conditional volatility in the forecasts of the covariance matrix required in asset allocation, risk management and asset pricing in the Vietnamese stock market. Also, emerging markets are very likely to exhibit characteristics different from those observed in developed markets. Emerging market returns and volatility are found to be more persistent than those in developed markets (Harvey, 1995, Bekaert and Harvey, 1997, Sadique and Silvapulle, 2001) . This could present some market inefficiency, or it could be because the risk factors are more persistent in emerging markets. Therefore, it should be of interests to study the long memory volatility behaviour in stock returns in Vietnam. Our paper thus fills in this gap, studying the economic benefits of employing multivariate short and long memory conditional volatility models to forecasts the covariance matrices in the Vietnamese stock market.

This paper evaluates the economic value of allowing for conditional volatility dynamics, short and long memory, in forecasting the covariance matrix for asset allocation over both short and long horizons, using the volatility timing framework of Fleming et al. (2001). Assuming constant expected returns, our investors follow a volatility timing strategy and periodically update their portfolios based on forecasts of the conditional covariance matrix. Dynamic portfolios constructed with alternative conditional volatility models are evaluated against static portfolios constructed with the constant unconditional covariance matrix estimates, and equally-weighted portfolios. We employ five multivariate conditional volatility models, both short memory and long memory. As is common in the literature, we restrict our attention to the class of EWMA and GARCH models. While many alternative multivariate conditional volatility models have been developed in literature, we choose models that can be parsimoniously constructed to forecast high-dimensional covariance matrices. The two short memory volatility models are the popularly used RiskMetrics EWMA model of JP Morgan (1994) and the GARCH model of Bollerslev (1986) embedded in the Dynamic Conditional Correlation (DCC) framework of Engle (2002) (the GARCH-DCC model). We also employ three long memory volatility models: the multivariate long memory EWMA model of Zumbach (2013) and two multivariate long memory models with the DCC structure. These are the component CGARCH model of Engle and Lee

(1999) and the FIGARCH model of Baillie et al. (1996). While the models that are based on the DCC decomposition allow for long memory only in the variances, the long memory EWMA model captures the long memory behaviour in both the variances and covariances.

We construct a high dimensional portfolio of VN30 components over the period 1 January 2010 to 30 June 2016. Portfolios generated using different conditional volatility models are evaluated using the out-of-sample Sharpe ratios and the performance fees that investors are willing to pay to switch from the static to the dynamic strategies. We also calculate the breakeven transaction costs that make investors indifferent between the static and the dynamic strategies in term of utility.

Our study is among the first to examine the economic value of multivariate conditional volatility models in the Vietnamese stock market. We report three main findings. First, consistent with the literature, the dynamic volatility timing strategies significantly outperform the static strategies with different performance measures and across different rebalancing frequencies. Second, there is negligible difference among short memory and long memory volatility models. However, due to their parsimony, the EWMA family models generally dominates the GARCH family. They consistently produce portfolios that are more economically useful than those produced by the GARCH volatility models at all investment horizons. Third, when transaction costs are taken into account, the gains from daily rebalanced dynamic portfolios deteriorate. However, it is still worth implementing the dynamic strategies at lower rebalancing frequencies. The results are robust to estimation error in expected returns, the choice of risk aversion coefficient and estimation windows.

The remaining of the paper is structured as follows. Section 2 provides details of the five multivariate conditional volatility models used in the empirical analysis. Section 3 sets up the asset allocation framework to study the economic usefulness of the dynamic strategies. Data is discussed in Section 4. Section 5 reports the empirical results, while Section 6 offers some concluding comments and suggestions for future research.

2 Multivariate Conditional Volatility Models

In the multivariate context, conditional volatility modelling poses significant computational challenges, especially for the high dimensional covariance matrices that are typically encountered in asset allocation and risk management. While many alternative volatility models have been developed in the literature, our choice reflects the need for parsimonious models that can be used for forecasting high dimensional covariance matrices. We employ two popularly

used multivariate short memory models: the GARCH(1,1)-DCC and the EWMA models. In order to evaluate the relative benefits of allowing for long memory when forecasting the covariance matrix, we compare the performance of the two short memory with three long memory volatility models: the multivariate long memory EWMA (LM-EWMA) model of Zumbach (2013) and the two DCC-structured long memory models: the FIGARCH(1, d ,1)-DCC and the CGARCH(1,1)-DCC models. In this section, we give details of each of these five models.

2.1 The GARCH(1,1)-DCC model

Observing that squared residuals are often autocorrelated even though residuals themselves are not, Engle (1982) sets the stage for the new class of time-varying conditional volatility models with the Autoregressive Conditional Heteroskedasticity (ARCH) model. Bollerslev (1986) extend the ARCH model by introducing autoregressive terms in his Generalised Autoregressive Conditional Heteroskedasticity (GARCH) model.

Consider a vector of log returns r_t with a conditional mean of zero and a conditional variance h_t :

$$r_t = \sqrt{h_t} \varepsilon_t \quad (1)$$

where ε_t denotes an i.i.d. mean zero, unit variance stochastic process. In the GARCH(1,1) model, the conditional variance, h_t , is modelled as

$$h_t = \omega + \alpha r_{t-1}^2 + \beta h_{t-1} \quad (2)$$

The parameter α determines the speed at which the conditional variance responds to new information, while the parameter $\alpha + \beta$ determines the speed at which the conditional variance reverts to its long run average. The GARCH model is a short memory model when the weights on past squared errors decline at an exponential rate. Assume that $\alpha + \beta < 1$ so that the long-run, or unconditional variance exists $\sigma^2 = \omega(1 - \alpha - \beta)^{-1}$, the h -step-ahead forecast of the GARCH(1,1) model is given by

$$h_{t+h} = \sigma^2 + (\alpha + \beta)^{h-1} (h_{t+1} - \sigma^2), \quad (3)$$

where σ^2 is the unconditional variance.

In order to implement the GARCH(1,1) model in the multivariate context, we employ the Dynamic Conditional Correlation (DCC) model of Engle (2002), in which the conditional covariance matrix is decomposed as follows:

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t \quad (4)$$

$$\mathbf{R}_t = \text{diag}\{\mathbf{Q}_t\}^{-\frac{1}{2}} \mathbf{Q}_t \text{diag}\{\mathbf{Q}_t\}^{-\frac{1}{2}} \quad (5)$$

$$\mathbf{Q}_t = \Omega + \delta \boldsymbol{\varepsilon}_{t-1} \boldsymbol{\varepsilon}_{t-1}' + \gamma \mathbf{Q}_{t-1} \quad (6)$$

where $\mathbf{R}_t$ is the conditional correlation matrix, $\mathbf{D}_t$ is a diagonal matrix with the time varying standard deviations $\sqrt{h_{i,t}}$ on the main diagonal, i.e., $\mathbf{D}_t = \text{diag}\{\sqrt{h_{i,t}}\}$, $\mathbf{Q}_t$ is the approximation of the conditional correlation matrix $\mathbf{R}_t$, and $\Omega = (1 - \delta - \gamma)\bar{\mathbf{R}}$, with $\bar{\mathbf{R}}$ being the unconditional average correlation $\bar{\mathbf{R}} = \frac{1}{T} \sum_{t=1}^T \boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t'$. The positive semi-definiteness of $\mathbf{Q}_t$ is guaranteed if δ and γ are positive with $\delta + \gamma < 1$ and the initial matrix $\mathbf{Q}_1$ is positive definite.

Here, we estimate the conditional volatility $\mathbf{D}_t$ by employing the GARCH model of Bollerslev (1986). We divide the returns by their conditional volatility and use the standardized, zero-mean residuals $\boldsymbol{\varepsilon}_t = \mathbf{D}_t^{-1} \mathbf{r}_t$ to compute the quasi-conditional correlation matrix $\mathbf{Q}_t$. As the diagonal elements of $\mathbf{Q}_t$ are equal to unity only on average, $\mathbf{Q}_t$ is rescaled to obtain the conditional correlation matrix $\mathbf{R}_t = \text{diag}\{\mathbf{Q}_t\}^{-\frac{1}{2}} \mathbf{Q}_t \text{diag}\{\mathbf{Q}_t\}^{-\frac{1}{2}}$. The conditional volatility $\mathbf{D}_t$ and conditional correlations $\mathbf{R}_t$ are then combined to estimate the conditional covariance matrix $\mathbf{H}_t$.

The h -step-ahead conditional covariance matrix is given by

$$\mathbf{H}_{t+h} = \mathbf{D}_{t+h} \mathbf{R}_{t+h} \mathbf{D}_{t+h}. \quad (7)$$

The forecast of each volatility in $\mathbf{D}_{t+h}$ is estimated using the recursive procedure, as in equation (3). Since $\mathbf{R}_t$ is a non-linear process, the h -step-ahead forecast of $\mathbf{R}_t$ cannot be computed using a recursive procedure. However, assuming for simplicity that $E_t(\boldsymbol{\varepsilon}_{t+1} \boldsymbol{\varepsilon}_{t+1}') \approx \mathbf{Q}_{t+1}$, Engle and Sheppard (2001) show that the forecasts of $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$ are given by

$$\mathbf{Q}_{t+h} = \sum_{j=0}^{h-2} (1-\delta-\lambda) \bar{\mathbf{Q}} (\delta+\gamma)^j + (\delta+\gamma)^{h-1} \mathbf{Q}_{t+1} \quad (8)$$

$$\mathbf{R}_{t+h} = \text{diag}\{\mathbf{Q}_{t+h}\}^{-\frac{1}{2}} \mathbf{Q}_{t+h} \text{diag}\{\mathbf{Q}_{t+h}\}^{-\frac{1}{2}}.$$

2.2 The RiskMetrics EWMA model

Consider an n -dimensional vector of returns $\mathbf{r}_t = (r_{1t}, r_{2t}, \dots, r_{nt})'$ with a conditional mean of zero and a conditional covariance matrix $\mathbf{H}_t$:

$$\mathbf{r}_t = \mathbf{H}_t^{\frac{1}{2}} \boldsymbol{\varepsilon}_t \quad (9)$$

where $\boldsymbol{\varepsilon}_t$ is i.i.d with $E(\boldsymbol{\varepsilon}_t) = 0$ and $\text{var}(\boldsymbol{\varepsilon}_t) = \mathbf{I}_n$. The short memory RiskMetrics EWMA covariance matrix is defined by

$$\mathbf{H}_t = \lambda \mathbf{H}_{t-1} + (1-\lambda) \mathbf{r}_{t-1} \mathbf{r}_{t-1}', \quad (10)$$

where λ is the decay factor $0 < \lambda < 1$. If the GARCH model normally assumes $\alpha + \beta < 1$, the EWMA model lets $\beta = \lambda$, $\alpha = 1 - \lambda$ and thus $\alpha + \beta = 1$. The EWMA model is a special case of the Integrated GARCH model of Engle and Bollerslev (1986) where a shock to volatility will eventually die out at an exponential rate, but it has a permanent effect on forecast volatility at all horizons. The EWMA model, though still a short memory model, can hence capture long run persistence in volatility. It is straightforward to show that the h -step cumulative forecast of the EWMA model is given by

$$\mathbf{H}_{t+1:t+h} = h \times \mathbf{H}_{t+1} \quad (11)$$

While the parameters of the GARCH family models have to be estimated by rigorous statistical methods, normally using the Maximum Likelihood procedure, the parameter λ of the EWMA process is often set *ad hoc*. As suggested by JP Morgan (1994), λ takes the values of 0.94 and 0.97 for daily and weekly forecasts, respectively.

2.3 The FIGARCH(1,d,1)-DCC model

Baillie et al. (1996) are among the first to develop long memory volatility models. Arguing that the volatility process is in a halfway house between $I(0)$ and $I(1)$, they propose the Fractionally Integrated GARCH (FIGARCH) model, in which long memory is introduced through a fractional difference operator, d . This model incorporates a slow hyperbolic decay for lagged squared innovations in the conditional variance, while still letting the cumulative impulse

response weights tend to zero, thus yielding a strictly stationary process. In the FIGARCH(1,d,1) model, the conditional volatility is modelled as:

$$h_t = \omega + [1 - \beta L - (1 - \phi L)(1 - L)^d] r_t^2 + \beta h_{t-1}, \quad (12)$$

with L being the lag operator. Baillie et al. (1996) show that for $0 < d \leq 1$, the FIGARCH process does not have a finite unconditional variance, and is not weakly stationary, a feature shared with the IGARCH model. However, by a direct extension of the corresponding proof for the IGARCH model, they show that the FIGARCH model is strictly stationary and ergodic. The FIGARCH process reduces to the GARCH process when $d = 0$. The h -step-ahead forecast of the FIGARCH(1,d,1) model is easily constructed by recursive substitution:

$$h_{t+h} = \omega(1 - \beta)^{-1} + [1 - (1 - \beta L)^{-1}(1 - \phi L)(1 - L)^d] h_{t+h-1}^2. \quad (13)$$

To implement the FIGARCH(1,d,1) model in the multivariate context, we use the DCC approach described above, with the same forecast functions for $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$.

2.4 The CGARCH(1,1)-DCC model

An alternative way to capture the long memory feature is through a component structure for volatility. Engle and Lee (1999) propose the Component GARCH (CGARCH) model, in which the long memory volatility process h_t is modelled as the sum of a long term trend component, q_t , and a short term transitory component, s_t . The CGARCH(1,1) model has the following specification:

$$h_t - q_t = \alpha(r_{t-1}^2 - q_{t-1}) + \beta(h_{t-1} - q_{t-1}) \quad (14)$$

$$q_t = \omega + \rho q_{t-1} + \phi(r_{t-1}^2 - h_{t-1}), \quad (15)$$

where $s_t = h_t - q_t$ is the transitory volatility component. The volatility innovation $r_{t-1}^2 - h_{t-1}$ drives both the trend and the transitory components. The long run component evolves over time following an AR process with ρ close to 1, while the short run component mean reverts to zero at a geometric rate $\alpha + \beta$. It is assumed that $0 < \alpha + \beta < \rho < 1$, meaning that the long run component is more persistent than the short run component. The h -step-ahead forecast of the CGARCH(1,1) model is given by

$$h_{t+h} = q_{t+h} + (\alpha + \beta)^{h-1} (h_t - q_t) \quad (16)$$

$$q_{t+h} = \frac{\omega}{1-\rho} + \rho^{h-1} \left(q_t - \frac{\omega}{1-\rho} \right). \quad (17)$$

As with the FIGARCH(1,d,1) model, in order to implement the CGARCH(1,1) model in the multivariate context, we use the DCC approach described above, with the same forecast functions for $\mathbf{Q}_{t+h}$ and $\mathbf{R}_{t+h}$.

2.5 The multivariate long memory LM-EWMA model

Zumbach (2013) generalises the EWMA model to incorporate long memory behaviour in volatility. The long memory LM-EWMA model is the weighted average of K standard (short memory) multivariate EWMA processes with logarithmically decaying weight:

$$w_k = \frac{1}{C} \left(1 - \frac{\ln(\tau_k)}{\ln(\tau_0)} \right) \text{ with the normalization constant } C = K - \sum_k \frac{\ln(\tau_k)}{\ln(\tau_0)}, \text{ such that } \sum_k w_k = 1.$$

The conditional covariance matrix in the LM-EWMA model can also be expressed as the weighted sum of the cross products of past returns:

$$\mathbf{H}_{t+1} = \sum_{i=0}^{\infty} \lambda(i) \mathbf{r}_{t-i} \mathbf{r}'_{t-i} \quad (18)$$

with $\sum \lambda(i) = 1$. In the RiskMetrics EWMA model of JP Morgan (1994), the weights $\lambda(i)$ decay geometrically, while in the LM-EWMA model, the weights $\lambda(i)$ are assumed to decay logarithmically, yielding a long memory process for the elements of the variance-covariance matrix:

$$\lambda(i) = \sum_k w_k (1 - \mu_k) \mu_k^i \quad (19)$$

where $\mu_k = \exp(-1/\tau_k)$, with geometric time structure $\tau_k = \tau_1 \sqrt{2}^{k-1}$ for $k = (1, \dots, K)$. The conditional covariance matrix is therefore defined parsimoniously as a process with just three parameters: τ_1 (the shortest time scale at which the volatility is measured, i.e. the lower cut-off), τ_K (the upper cut-off, which increases exponentially with the number of components K), and τ_0 (the logarithmic decay factor). For the univariate case, Zumbach (2006) sets the optimal parameter values at $\tau_0 = 1560$ days = 6 years, $\tau_1 = 4$ days and $\tau_K = 512$ days, which is equivalent to $K = 15$.

Since the LM-EWMA covariance matrix is the sum of EWMA processes over increasing time horizons, forecasts of the covariance matrix are straightforward to obtain using a recursive procedure (see Zumbach, 2006, for details of the univariate case). The 1-step-ahead forecast of the covariance matrix is already given by Equation (18). Under the assumption of serially uncorrelated returns, the h -step-ahead cumulative forecast of the covariance matrix given the information set F_t at time t is equal to:

$$\mathbf{H}_{t+1:t+h} = h \sum_{i=0}^T \lambda(h, i) \mathbf{r}_{t-i} \mathbf{r}'_{t-i} \quad (20)$$

with the weights $\lambda(h, i)$ being given by

$$\lambda(h, i) = \sum_{k=1}^K \frac{1}{h} \sum_{j=1}^{h-1} w_{j,k} \frac{(1-\mu_k)}{1-\mu_k^T} \mu_k^i, \quad (21)$$

where T is the cut-off time², $w_{j,k}$ is the k^{th} element of vector $\mathbf{w}_j = \mathbf{w}'(\mathbf{M} + (\mathbf{I} - \boldsymbol{\mu})\mathbf{w}')^j$, $\mathbf{w} = (w_1, w_2, \dots, w_K)'$, $\boldsymbol{\mu}$ is the vector of μ_k , $\mathbf{M}$ is the diagonal matrix consisting of μ_k , and $\mathbf{I}$ is the unit vector. Since $\sum_k w_k = 1$, we obtain $\sum \lambda(h, i) = 1$.

Note that when $K=1$, we have $w=1$, and the LM-EWMA forecast function reduces to a standard short memory EWMA forecast function with forecast weights $\lambda(h, i) = (1-\mu_k) \mu_k^i / (1-\mu_k^T)$, independent of the forecast horizon. Since the weights $\lambda(h, i)$ are set *ad hoc*, the forecast in equation is straightforward to compute. As with the standard EWMA model, the LM-EWMA model circumvents the computational burden of other multivariate long memory models, and indeed, can be implemented in a spreadsheet easily.

3 The Economic Value of Dynamic Volatility Timing Strategy

3.1 Dynamic Volatility Timing Framework

We consider a risk averse investor who wants to maximise his expected utility U_{t+1} in the mean-variance optimization framework. He will allocate a fraction $\mathbf{w}_t$ of his wealth to n risky assets and the remainder $(1-\mathbf{w}_t' \mathbf{1})$ to a risk-free asset, where $\mathbf{1}$ is the $n \times 1$ unit vector, so that:

$$\max_{\mathbf{w}_t} \left\{ E(U_{t+1}) = \mu_{p,t+1} - \frac{\lambda}{2} \sigma_{p,t+1}^2 \right\} \quad (22)$$

² Zumbach (2011) suggests that for many practical applications, the memory length T is of the order of one to two years ($T = 260$ to $T = 520$). Here, we choose T equal to the estimation window length.

where $\mu_{p,t+1}$ is the portfolio's expected returns $\mu_{p,t+1} = \mathbf{w}_t' \boldsymbol{\mu}_{t+1} + (1 - \mathbf{w}_t' \mathbf{1}) r^f$, $\sigma_{p,t+1}^2$ is the portfolio's expected variance $\sigma_{p,t+1}^2 = \mathbf{w}_t' \mathbf{H}_{t+1} \mathbf{w}_t$, $\boldsymbol{\mu}_{t+1}$ is the vector of expected returns, $\mathbf{H}_{t+1}$ is the conditional covariance matrix, r^f is the risk-free rate and λ is the risk aversion coefficient. In the empirical study, we assume a risk free rate of 7% and a risk aversion coefficient of 1. Different values of λ are later considered in the robustness test. Short sales are allowed and no transaction costs are included. The solution to this optimization problem is:

$$\mathbf{w}_t^* = \frac{1}{\lambda} \mathbf{H}_{t+1}^{-1} (\boldsymbol{\mu}_{t+1} - r^f \mathbf{1}). \quad (23)$$

Following Fleming et al. (2001), the investor models expected returns as constant ($\boldsymbol{\mu}_{t+1} \equiv \boldsymbol{\mu}$). If the investor also assumes a constant covariance matrix ($\mathbf{H}_{t+1} \equiv \mathbf{H}$), the optimal weights will be constant over time and he follows a '*static strategy*'. However, if the investor believes that the covariance matrix is time-varying, he will follow a '*dynamic strategy*' to change the optimal weights based on his forecasts of the conditional covariance matrix. The investor will employ the five multivariate conditional volatility models (the two short memory multivariate EWMA and the GARCH(1,1)-DCC models and the three long memory LM-EWMA, FIGARCH(1,d,1)-DCC and CGARCH(1,1)-DCC models) to generate forecasts of the covariance matrix for the dynamic strategies. The economic value of volatility timing can be evaluated by comparing the performance of the static and dynamic portfolios. The portfolios constructed with the three multivariate long memory volatility models are also compared to those constructed with the two short memory volatility models to specifically evaluate the gains of exploiting long memory vs. short memory properties of volatility.

3.2 Performance Measures of Dynamic Strategies

The performance of the optimal portfolios is evaluated using two common measures. First, we estimate the out of sample Sharpe ratio, the most commonly used performance measure in literature. The Sharpe ratio for each strategy is calculated as the sample mean of the realised portfolio excess returns over the risk free rate divided by their sample standard deviations,

$$SR = \frac{(\mu_p - r_f)}{\sigma_p}. \quad \text{Second, following Fleming et al. (2001), we use a utility-based approach to}$$

measure the value of the performance gains associated with using a given estimator of the conditional covariance matrix. In so doing, we estimate the performance fee, Δ , defined as the maximum fee that the investor would be willing to pay to switch from a static strategy to a

dynamic strategy, without being worse off in terms of utility. To estimate this fee, we find the value of Δ that equates the realised average utilities for two alternative portfolios:

$$\sum_{t=0}^{T-1} (R_{d,t+1} - \Delta) - \frac{\gamma}{2(1+\gamma)} (R_{d,t+1} - \Delta)^2 = \sum_{t=0}^{T-1} R_{s,t+1} - \frac{\gamma}{2(1+\gamma)} R_{s,t+1}^2, \quad (24)$$

where γ_t is the coefficient of relative risk aversion, $R_{d,t+1}$ and $R_{s,t+1}$ are the gross realised returns of the dynamic and static strategies, respectively. In the empirical analysis, we report the annualised performance fees in basis points for two different values of γ , 1 and 5.

3.3 Transaction Costs

Volatility timing requires regular updates of portfolios, thus incurring non-trivial transaction costs. Transaction costs may be high enough to offset all the gains that arise from the dynamic strategy. Following Han (2006), we estimate the breakeven transaction cost τ^{be} , defined as the transaction cost that make investors indifferent between the dynamic and the static strategies in terms of utility. If an investor has a transaction cost lower than the breakeven transaction cost, he will be better off with the dynamic strategy; otherwise he should follow the static benchmark. Han sets the transaction costs equal to a fixed percentage (τ) of the value traded for all stocks. The costs for the static and dynamic strategies are given by

$$\tau \left| w - \frac{w(1+r_{p,t+1})}{w(r_{p,t+1} - r_f) + r_f + 1} \right| \quad (25)$$

$$\text{and } \tau \left| w_{t+1} - \frac{w_t(1+r_{p,t+1})}{w_t(r_{p,t+1} - r_f) + r_f + 1} \right|, \text{ respectively.} \quad (26)$$

The breakeven transaction cost is computed by equating the utilities of the static and dynamic strategies in Equation (25) after taking into account the trading costs. The higher the breakeven transaction cost, the more easily the dynamic trading strategies can be implemented. Since the breakeven transaction is a proportional cost paid every time the portfolios are rebalanced, we report this cost in basis points at the rebalancing frequency, e.g., for a daily rebalanced portfolio, we report the cost in daily basis points. Breakeven transaction costs are only estimated when the performance fees in Equation (24) are positive.

4 Data Description

We construct a high dimensional portfolio, comprising the components of the VN30 index as of 30 June 2016. Daily data are collected from FiinPro for the period from 01 January 2010 to 30 June 2016. As the components of our portfolio remain unchanged during our experiments, we exclude four stocks (MBB, HHS, NT2 and FLC), which were listed after January 2010. Returns are calculated as the log price difference over consecutive days. All days on which the market was closed are excluded from the sample, yielding 1613 observations. Summary statistics for the 26 VN30 stocks are given in Table 1.

Table 1. Summary Statistics for the VN30 Portfolio

Return series	Mean (%)	Std. Dev. (%)	Skewness	Kurtosis	JB	Long memory test		
						d _{GPH}	R/S	V/S
BVH	13.12	44.52	0.09	2.74	6.67	0.35***	1.92***	0.23
CII	5.37	36.60	0.08	3.67	32.19	0.38***	1.93***	0.28***
CTG	3.86	31.00	0.24	4.31	131.11	0.34***	2.03***	0.31***
DPM	6.68	28.42	-0.05	4.40	131.62	0.30***	2.41***	0.53***
EIB	-1.42	24.83	0.26	6.08	657.59	0.23***	1.53	0.12
FPT	9.06	26.17	0.15	5.13	312.05	0.20**	1.45	0.11
GMD	-3.35	37.11	0.04	3.64	27.59	0.19**	2.33***	0.58***
HAG	-25.16	35.54	0.11	3.85	52.32	0.41***	1.93***	0.27***
HCM	2.77	38.42	0.09	3.35	10.57	0.43***	1.38	0.13
HPG	15.15	33.77	0.14	3.69	37.25	0.09	1.75**	0.24***
HSG	18.16	41.08	0.08	3.29	7.36	0.17*	2.00***	0.32***
HVG	-2.11	39.42	0.18	3.53	27.85	-0.15	0.86	0.03
ITA	-20.37	43.27	0.16	2.79	10.00	0.36***	2.33***	0.54***
KBC	-14.81	45.77	0.07	2.59	12.68	0.28***	2.85***	0.86***
KDC	8.67	30.30	0.16	4.64	186.61	0.12	1.51	0.11
MSN	9.55	34.04	0.09	3.67	32.46	0.49***	2.38***	0.55***
PPC	1.59	38.05	0.03	3.46	14.51	0.24***	2.23***	0.35***
PVD	-5.39	36.40	0.04	3.74	37.27	0.14*	1.35	0.09
PVT	-0.77	44.05	0.16	2.85	8.02	0.04	2.01***	0.36***
REE	8.30	31.53	0.14	4.05	78.78	0.26***	2.16***	0.33***
SBT	25.81	33.80	0.16	3.91	63.32	0.38***	2.04***	0.31***
SSI	-3.93	34.67	0.18	3.61	33.48	0.48***	2.12***	0.35***
STB	-1.19	28.86	0.30	4.89	264.63	0.23***	1.02	0.08
VCB	9.88	32.15	0.06	4.12	84.97	0.18*	1.34	0.11
VIC	18.89	30.44	0.07	4.02	71.03	0.32***	2.70***	0.63***
VNM	30.74	23.63	0.26	5.63	484.33	0.21***	1.12	0.07

*The table reports descriptive statistics for the daily returns on 26 components of the VN30 index. Means and standard deviations are annualised. The sample period is from 01 January 2010 to 30 June 2016. d_{GPH} is the fractional difference operator estimated using the Geweke-Porter-Hudak (GPH) test. R/S and V/S report the statistics to detect for long memory using the rescaled range estimator of Lo (1991) and of Giraitis et al. (2003), respectively. Rejection of the null hypothesis $H_0: d = 0$ (short memory) is displayed by *, ** and *** for 10%, 5% and 1% significance level.*

The performance of the 26 VN30 stocks differ dramatically. While VNM makes an average annual return of 30.74% over the period, HAG loses, on average, 25.16% per year. These stocks are characterised by a high level of volatility. However, higher volatility does not necessarily come with higher returns. VNM has the highest returns with the lowest volatility, while some other stocks such as ITA, KBC, HAG have negative returns with much higher volatility. The return series are highly non-normal, with a high leptokurtosis. The average correlation coefficient of the VN30 components is 0.34.

We conduct tests to confirm the evidence of long memory dynamics in the volatility. A series is said to have a long memory if its fractional difference parameter, $d \in (0,1]$. We apply three extensively used tests of long memory, namely the semiparametric (GPH) estimator of Geweke and Porter-Hudak (1983), the nonparametric modified ‘rescaled range’ (R/S) test of Lo (1991) and the rescaled variance (V/S) test of Giraitis et al. (2003). The GPH estimator report the fractional difference operators d , while the two modified R/S and V/S tests report the statistics to test for the null hypothesis $H_0: d = 0$ (short memory) against $H_1: d > 0$. To conduct the GPH test, we use the recommended bandwidths equal to the square root of the sample size, $m = 40$. For the R/S and V/S test, we choose $q = 21$ to include autocovariances of months. Most of the return series show evidence of long memory in the volatility in at least one of the tests, with the only exception of KDC and HVG. Most of the fractional difference parameters in the GPH tests are significantly greater than zero. We also conduct a one-sided test of the hypothesis $d = 0.5$, against the alternative $d < 0.5$. Rejecting this hypothesis, we confirm that the volatility processes of these series are characterised by long memory, but are nevertheless stationary.

[Insert Table 1 here]

5 Empirical Results

5.1 Performance Analysis of the Dynamic Asset Allocation Strategies

The whole sample is divided into an estimation period and a forecast period. The estimation period is from 1 Jan 2010 to 31 Dec 2013 (996 observations) and the forecast period from 1 Jan 2014 to 30 June 2016 (617 observations). Expected returns are assumed to be constant and be the sample mean of the estimation period. The investor actively rebalances his portfolios periodically, based on changes in the forecasts of the conditional covariance matrix. The estimation period is used to initiate the estimation of the conditional covariance matrices and generate one step ahead forecasts. The forecasts are then used to compute the optimal portfolio weights. Realised portfolio returns at the next step are calculated. Then the estimation window is

rolled forward one step, models re-estimated, forecasts made, and portfolios rebalanced until the end of the sample is reached. The realised performance of the dynamic portfolios will be compared with that of the ex ante optimal static portfolio, constructed based on the sample mean and covariance matrix of the estimation period. Another benchmark is the equally weighted portfolio.

Table 2 evaluates the out of sample performance of the daily rebalanced VN30 portfolios. The VN30 portfolio is constructed from 26 components of the VN30 Index. Investors would generally be better off with the dynamic strategies than the static and the equally weighted strategies. A day trader, for example, would be willing to pay up to 288 bps to switch from the static to the dynamic LM-EWMA strategy. However, in terms of the Sharpe ratio, some of the dynamic strategies (the GARCH-DCC, CGARCH-DCC and FIGARCH-DCC models) fail to dominate the static portfolios. Among the dynamic strategies, the EWMA and LM-EWMA models consistently outperform across all performance measures. With high breakeven transaction costs, it is also more feasible to implement the EWMA and LM-EWMA models in practice. It is interesting to see dominance of the portfolio constructed with the EWMA models. The outperformance of the EWMA model may be attributed to its implementation simplicity, hence yielding less estimation error. The EWMA processes are also not constrained by a mean level of volatility as in the GARCH family model and thus can be readily adjusted to changes in unconditional volatility.

Table 2. Portfolio Performance of the VN30 Portfolio

	μ (%)	σ (%)	SR	Δ_1	Δ_5	τ_1	τ_5
1/N	8.178	19.070	0.062				
Static	8.586	2.680	0.592				
<i>Volatility timing strategies</i>							
EWMA	25.510	14.242	1.300	1595	1203	14	11
GARCH-DCC	9.180	4.751	0.459	52	21	1	1
LM-EWMA	11.597	5.841	0.787	288	234	8	7
CGARCH-DCC	7.936	4.249	0.220	-70	-92	—	—
FIGARCH-DCC	9.940	7.657	0.384	105	-1	3	—

The table compares the out-of-sample performance of the optimal VN30 portfolio. The static portfolio is constructed using the constant mean and covariance matrix of the estimation period. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the annualised performance fee (in basis points) Δ_γ that an investor with a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the breakeven transaction cost τ_γ (in basis points) that he will be better off with the dynamic strategy.

The results in Table 2 suggests that volatility timing can generate economic value. However, the mean-variance portfolio optimisation requires the forecasts of both expected returns and covariance matrix. Expected returns here are assumed to the sample mean of the estimation period so that we can consider the benefits of conditional volatility estimators. However, portfolio performance is highly sensitive to the expected returns and considering only one set of expected returns may not appropriate. To assess whether the economic value of volatility timing strategies is realizable we need to incorporate the effects of estimation errors in expected returns.

5.2 Controlling for Estimation Error in Expected Returns and Longer Horizon Forecasts

To account for estimation error in expected returns, we follow Fleming et al.'s (2001) recommendation to consider a range of expected returns that are generated via a bootstrap procedure. An artificial sample of 2,000 observations is created by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. We then estimate the unconditional mean and covariance matrix of this artificial return series. Dynamic portfolios are constructed using the constant unconditional expected returns from the bootstrap and forecasts of the conditional covariance matrix. To ensure the static and the dynamic portfolios are based on the same ex ante information, the static benchmark portfolio is formed using the bootstrap constant expected returns and covariance matrix. We repeat this procedure with 1000 trials, studying the economic gains of volatility timing across a wide range of plausible vectors of expected returns. We also study the benefits of the dynamic strategies with longer weekly and monthly investment horizons.

Table 3 summarises the average results across the 1,000 bootstrap vectors of expected returns for the VN30 portfolio. Consistent with the literature ((see, for example, Fleming et al., 2001, Fleming et al., 2003, Han, 2006, Hyde et al., 2010), the dynamic strategies consistently outperform the static with all performance measures and rebalancing frequencies. The dynamic portfolios are generally riskier than the static portfolios, but they generate much higher returns, hence yielding high Sharpe ratios. The Sharpe ratios of the dynamics portfolios are around 2 times as much as those of the static portfolios. The daily rebalanced LM-EWMA portfolio, for example, outperforms the static portfolio in terms of the Sharpe ratio in 97.5% of all bootstrap vectors of returns. The dynamic strategies also yield large performance fees due to their high realised returns. The investor is also better off by at least 205 *bps* and up to 2446 *bps* annually, on average, when switching to the dynamic portfolios. It is interesting to see that moving from daily to lower weekly and monthly rebalancing frequencies reduces the Sharpe ratio and the performance fees, suggesting the benefits of high frequency trading in Vietnam. Once we

consider transaction costs, however, however, the dynamic strategies are only attractive for lower rebalancing frequencies. For example, a daily trader with $\gamma=1$ is only better off with the LM-EWMA portfolio if his realised transaction cost is lower than 18 *bps*. With the average transaction costs of 20-30 *bps* currently in Vietnam, it may be infeasible to rebalance portfolios every day. The breakeven transaction costs of a weekly and monthly trader are much higher than those of a daily trader due to less frequent rebalancing. It is evident that the conditional volatility portfolios are more feasible in terms of transaction costs with low frequent trading.

Among the conditional volatility models, the EWMA and LM-EWMA models outperform significantly. The LM-EWMA model tends to outperform the EWMA models in terms of the Sharpe ratio with longer forecast horizons, however, the EWMA model still dominates in terms of performance fee due to high realised returns. There is negligible difference between short memory and long memory volatility models. Though most of the volatility series exhibit long memory behaviour, the long memory volatility models fail to capture this feature and produce superior forecasts. It may be because the long memory models may be specified correctly, but their complex structure may hinder their performance. The high level of parameterization of the Component GARCH and FIGARCH models evidently generates large estimation errors that are detrimental to their performance. The long memory CGARCH-DCC and FIGARCH-DCC models in general cannot outperform the short memory EWMA, GARCH-DCC and the much simpler long memory LM-EWMA models. Parsimony may also explain the outperformance of the EWMA family models as compared to the GARCH family models. In particular, the simplicity in estimation of the two EWMA and LM-EWMA models is evidently beneficial in the high dimensional case.

Table 3. Average Portfolio Performance of the VN30 Portfolio with Bootstrap Experiments

	μ (%)	σ (%)	SR	<i>p-value</i>	Δ_1	Δ_5	τ_1	τ_5
Panel A. Daily rebalancing								
Static	9.976	4.604	0.668					
<i>Volatility timing strategies</i>								
EWMA	36.027	18.316	1.584	0.967	2446	1807	17	12
GARCH-DCC	15.196	6.098	1.341	0.942	514	482	10	10
LM-EWMA	18.522	7.774	1.479	0.975	835	756	18	16
CGARCH-DCC	13.672	5.519	1.209	0.881	365	347	9	9
FIGARCH-DCC	17.342	8.978	1.152	0.782	826	644	5	4
Panel B. Weekly rebalancing								
Static	10.177	5.007	0.661					
<i>Volatility timing strategies</i>								
EWMA	19.083	8.333	1.452	0.955	869	779	40	36

GARCH-DCC	13.942	5.144	1.348	0.948	376	375	31	31
LM-EWMA	17.345	6.949	1.489	0.966	705	659	34	32
CGARCH-DCC	14.123	5.083	1.399	0.959	395	394	38	38
FIGARCH-DCC	18.864	9.658	1.233	0.893	834	695	19	16
Panel C. Monthly rebalancing								
Static	9.896	6.614	0.458					
<i>Volatility timing strategies</i>								
EWMA	17.711	11.236	0.961	0.896	741	569	68	52
GARCH-DCC	12.583	4.998	1.136	0.972	278	319	123	140
LM-EWMA	15.317	7.587	1.101	0.960	536	509	85	81
CGARCH-DCC	11.924	6.262	0.848	0.800	205	214	71	74
FIGARCH-DCC	16.397	9.213	1.050	0.910	630	544	67	58

The table reports the average out-of-sample performance of the VN30 portfolio across a wide range of bootstrap-generated expected returns. An artificial sample of 2,000 observations is generated by randomly picking up blocks, with replacement, of 15 observations from the series of actual returns. The procedure is repeated with 1,000 trials. Panels A, B and C report the results of the daily, weekly and monthly rebalanced portfolios, respectively. The static portfolios are constructed using the bootstrap expected returns and covariance matrices. For each dynamic strategy, the table reports the average annualised realised return (μ), the average annualised realised volatility (σ), the average Sharpe ratio (SR), the p-value (proportion) that the dynamic strategy outperforms the static alternative in terms of the Sharpe ratio, the average annualised performance fee (in basis points) Δ_γ that an investor a constant relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the average breakeven transaction cost τ_γ (in basis points) that he will be better off with the dynamic strategy.

5.3 Sensitivity to Risk Aversion Coefficient

In this section, we evaluate the performance of the dynamic strategies, controlling for different risk aversion coefficients λ . So far all reported results are based on $\lambda=1$. For each value of the risk aversion coefficients λ , we again generate 1,000 bootstrap vectors of expected returns and use them, along with the conditional covariance matrix estimates, to construct the optimal portfolios. Table 4 evaluate the performance of the dynamic long memory and short memory volatility timing strategies against the static strategies. To save space, we only report the performance of the long memory LM-EWMA and the short memory EWMA models. These are the two best performing models in the previous experiments.

Not surprisingly, when the investor is more risk averse, he will choose portfolios with lower risk, accepting lower expected returns, and paying lower performance fees. The Sharpe ratios are approximately the same for all risk aversion coefficients, with the slight difference due to the bootstrap procedure. Again, the dynamic strategies – both long memory and short memory- consistently dominate the static strategies in both datasets with all rebalancing frequencies and

risk aversion coefficients. The long memory volatility timing portfolios, for example, yield higher Sharpe ratios than the static portfolios in over 95% of total trials. With high breakeven transaction costs, the dynamic strategies are feasible in longer forecast horizons, where the LM-EWMA portfolios generally yield higher Sharpe ratios than the short memory EWMA portfolios. However, the LM-EWMA model underperform their short memory counterparts in terms of performance fees due to lower realised returns.

Table 4. Comparison of the Static and the Dynamic Volatility Timing Strategies Using Different Risk Aversion Coefficients

The table compares the average out-of-sample performance of the static and dynamic strategies using different risk aversion coefficients λ . A bootstrap procedure is applied to control for estimation error in expected returns. The static portfolios are constructed using the bootstrap expected returns and covariance matrices. The short memory portfolios are constructed with the EWMA model, while the long memory portfolios are constructed with the LM-EWMA model. For each dynamic strategy, the table reports the average annualised realised return (μ), the annualised realised volatility (σ), the Sharpe ratio (SR), the p -value (proportion) that the dynamic strategy outperforms the static alternative in terms of the Sharpe ratio, the annualised performance fee (in basis points) Δ_γ that an investor with a relative risk coefficient of γ is willing to pay to switch from the static portfolio to the dynamic portfolio, and the average breakeven transaction cost τ_γ (in basis points) that he will be better off with the dynamic strategy.

	<i>Static</i>			<i>Short memory</i>			<i>Short memory vs. Static</i>			<i>Long memory</i>			<i>Long memory vs. Static</i>		
λ	μ (%)	σ (%)	SR	μ (%)	σ (%)	SR	p -value	Δ_1	τ_1	μ (%)	σ (%)	SR	p -value	Δ_1	τ_1
Panel A. Daily rebalancing															
1	9.976	4.604	0.668	36.027	18.316	1.584	0.967	2446	17	18.522	7.774	1.479	0.975	835	18
2	8.488	2.302	0.668	21.514	9.158	1.584	0.967	1263	17	12.761	3.887	1.479	0.975	422	18
3	8.018	1.516	0.696	16.877	6.103	1.616	0.967	868	18	10.969	2.592	1.528	0.981	293	19
4	7.762	1.145	0.688	14.473	4.568	1.632	0.976	661	18	9.987	1.940	1.535	0.979	221	19
5	7.613	0.927	0.683	13.015	3.670	1.636	0.974	534	18	9.399	1.559	1.537	0.976	178	19
Panel B. Weekly rebalancing															
1	10.177	5.007	0.661	19.083	8.333	1.452	0.955	869	40	17.345	6.949	1.489	0.966	705	34
2	8.586	2.563	0.643	12.934	4.169	1.426	0.943	429	39	12.116	3.483	1.470	0.963	350	34
3	8.055	1.664	0.665	10.940	2.753	1.437	0.938	286	40	10.392	2.294	1.483	0.955	232	34
4	7.788	1.254	0.658	9.978	2.096	1.426	0.995	218	40	9.553	1.744	1.467	0.963	176	34
5	7.626	1.010	0.642	9.382	1.660	1.439	0.942	175	40	9.041	1.386	1.477	0.950	141	34
Panel C. Monthly rebalancing															
1	9.896	6.614	0.458	17.711	11.236	0.961	0.896	741	68	15.317	7.587	1.101	0.960	536	85
2	8.420	3.302	0.452	12.165	5.646	0.927	0.881	364	67	11.033	3.799	1.069	0.954	260	83
3	7.956	2.182	0.462	10.483	3.758	0.938	0.889	248	69	9.720	2.534	1.078	0.951	176	84
4	7.706	1.634	0.452	9.547	2.826	0.919	0.793	181	68	9.010	1.901	1.068	0.951	130	83
5	7.568	1.317	0.456	9.089	2.259	0.938	0.886	150	70	8.627	1.518	1.079	0.957	106	85

5.4 Sensitivity to the Estimation Windows

As the performance of conditional volatility models may be sensitive to the sample length used in their estimation, we investigate the performance of the strategies using a range of estimation windows. In particular, we consider estimation windows of 2, 3, 4 and 5 years of daily data. The analysis is again conducted with the bootstrap vectors of expected returns. We exclude the FIGARCH-DCC model as its estimation requires a prohibitively high upper cut-off point (the truncation lag for the FIGARCH model is normally set at 1000). To save space, Figure 1 only reports the average Sharpe ratios of the daily and weekly rebalanced dynamic portfolios using different estimation windows. As with the previous results, the dynamic strategies outperform the static strategy across different estimation windows. The Sharpe ratios of the dynamic strategies is around three times those of the static strategies. The dynamic strategies also yield large performance fees due to their high realised returns. Among the dynamic portfolios, the portfolios constructed from the more parsimonious EWMA and LMEWMA models, again, consistently dominate the more complex GARCH and CGARCH portfolios. It is notable that the Sharpe ratios of the conditional volatility models tend to increase with the estimation window length, suggesting that the more observations in the estimation period, the less noisy the estimates and the more accurate the forecast of the conditional covariance matrix.




Figure 1. Sensitivity to Estimation Windows: Sharpe Ratios of the Dynamic Portfolios.

The average Sharpe ratios of the optimal portfolios constructed from different volatility models are estimated with different estimation windows. Bootstrapped expected returns are employed to account for estimation error. The estimation windows correspond to 2, 3, 4 and 5 years of daily data.

6 Conclusion

The paper examines the economic value of allowing for conditional volatility in forecasting the conditional covariance matrix for dynamic asset allocation. Consistent with the literature, the results clearly demonstrate that investors are willing to pay to switch from the static unconditional strategy to the dynamic volatility timing alternatives. However, the findings show little difference among the forecast performance between short and memory volatility models. Among the conditional volatility models, the more parsimonious EWMA and LM-EWMA models nevertheless dominate. They consistently produce portfolios that have more economic value than those produced by more complex GARCH-type with all performance measures and across all investment horizons. The high degree of parameterisation and computational burden may generate such high estimation error that it is detrimental to the performance of the GARCH models. When transaction costs are considered, however, the dynamic volatility timing strategies are only attractive at lower rebalancing frequencies. The results are robust to estimation error in expected returns, the choice of risk aversion coefficient and estimation windows.

The analysis of the economic value of multivariate conditional volatility models can also be extended in several ways. First, we assume constant expected returns. Time-varying volatility affects returns and it is hard to justify the separation of the dynamics of expected returns from those of volatility and correlation. It would thus be of interest to extend the study in the context of time-varying expected returns. Second, it may be useful to study the economic benefits of dynamic strategies in an intertemporal asset allocation framework, since dynamic strategies may behave differently in the presence of hedging demands. Third, the study is restricted to examining the economic value of forecasting the conditional covariance matrix from an investment perspective. One may want to examine the implications for the conditional covariance matrix in other situations encountered in practice, such as risk management.

References

- BAILLIE, R. T., BOLLERSLEV, T. & MIKKELSEN, H. O. 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74, 3--30.
- BEKAERT, G. & HARVEY, C. R. 1997. Emerging equity market volatility. *Journal of Financial economics*, 43, 29-77.

- BOLLERSLEV, T. 1986. Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31, 307--327.
- DAVIDSON, J. 2004. Moment and memory properties of linear conditional heteroscedasticity models, and a new model. *Journal of Business & Economic Statistics*, 22, 16--29.
- ENGLE, R. F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50, 987--1007.
- ENGLE, R. F. 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20, 339--350.
- ENGLE, R. F. & BOLLERSLEV, T. 1986. Modelling the persistence of conditional variances. *Econometric Reviews*, 5, 1--50.
- ENGLE, R. F. & COLACITO, R. 2006. Testing and valuing dynamic correlations for asset allocation. *Journal of Business and Economic Statistics*, 24, 238--253.
- ENGLE, R. F. & LEE, G. 1999. A Long Run and Short Run Component Model of Stock Return Volatility In: ENGLE, R. & WHITE, H. (eds.) *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive WJ Granger*. New York: Oxford University Press.
- ENGLE, R. F. & SHEPPARD, K. 2001. Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. *NBER Working Paper*.
- FLEMING, J., KIRBY, C. & OSTDIEK, B. 2001. The economic value of volatility timing. *The Journal of Finance*, 56, 329--352.
- FLEMING, J., KIRBY, C. & OSTDIEK, B. 2003. The economic value of volatility timing using "realized" volatility. *Journal of Financial Economics*, 67, 473--509.
- GEWEKE, J. & PORTER-HUDAK, S. 1983. The estimation and application of long memory time series models. *Journal of Time Series Analysis*, 4, 221--238.
- GIRAITIS, L., KOKOSZKA, P., LEIPUS, R. & TEYSSIÈRE, G. 2003. Rescaled variance and related tests for long memory in volatility and levels. *Journal of Econometrics*, 112, 265-294.
- HAN, Y. 2006. Asset allocation with a high dimensional latent factor stochastic volatility model. *The Review of Financial Studies*, 19, 237--271.
- HARRIS, R. D. F. & NGUYEN, A. 2013. Long Memory Conditional Volatility and Dynamic Asset Allocation. *International Journal of Forecasting*, 29, 258--273.
- HARVEY, C. R. 1995. Predictable risk and returns in emerging markets. *Review of Financial studies*, 8, 773-816.
- HYDE, S., NGUYEN, N. & POON, S.-H. 2010. International Portfolio Choice and the Economic Value of Jump Timing. *Seminars Inquire Europe*. Rome, Italy.
- JPMORGAN 1994. RiskMetrics™, Technical Document. JP Morgan.
- LO, A. W. 1991. Long-Term Memory in Stock Market Prices. *Econometrica*, 59, 1279-1313.
- NIGUEZ, T. M. & RUBIA, A. 2006. Forecasting the conditional covariance matrix of a portfolio under long-run temporal dependence. *Journal of Forecasting*, 25, 439--458.
- SADIQUE, S. & SILVAPULLE, P. 2001. Long-term memory in stock market returns: international evidence. *International Journal of Finance & Economics*, 6, 59-67.
- TEYSSIERE, G. 1998. Multivariate long-memory ARCH modelling for high frequency foreign exchange rates. Citeseer.
- THORP, S. & MILUNOVICH, G. 2007. Symmetric versus asymmetric conditional covariance forecasts: Does it pay to switch? *Journal of Financial Research*, 30, 355--377.
- TRAN MANH TUYEN 2011. Modeling Volatility Using GARCH Models: Evidence from Vietnam. *Economic Bulletin*, 31, 1935--1942.
- VO XUAN VINH & NGUYEN THI KIM NGAN 2011. Volatility in Stock Return Series of Vietnam Stock Market. *Science & Technology Development*, , 14, 5--21.
- ZUMBACH, G. 2006. The Riskmetrics 2006 Methodology. *Technical Report*, RiskMetrics Group.

ZUMBACH, G. 2013. The statistical properties of the innovations in multivariate ARCH processes in high dimensions. *Quantitative Finance*, 13, 29-44.