The Limits to Minimum-Variance Hedging

Richard D. F. Harris, *University of Exeter*
Jian Shen, *University of Exeter*
Evarist Stoja, *Queen's University of Belfast*

Paper Number: 07/12

October 2007

Abstract

Models of the time-varying conditional minimum-variance hedge ratio (MVHR) typically do not provide a significant improvement in terms of hedging performance over the unconditional MVHR model. In view of the widely documented success of conditional volatility models (on which models of the conditional MVHR are usually based), this is somewhat surprising. In this paper, using the recently developed realized beta framework of Andersen, Bollerslev, Diebold and Wu (2005), we explore the reasons for this finding. We firstly show that the reduction in hedged portfolio variance that conditional MVHR models offer falls far short of the *ex post* maximal reduction in variance obtained using an estimate of the unobserved 'integrated' MVHR. We investigate the statistical properties of the forecasts of conditional MVHR models and show that while they do contain significant information about the integrated MVHR, they are systematically biased and inefficient. However, correcting for this bias and inefficiency does little to improve their hedging performance, suggesting that their poor performance is more likely to be attributable to the unpredictability of the integrated MVHR.

Keywords: Minimum-variance hedge ratio; Realized beta; Multivariate conditional volatility models; Bias correction.

Address for correspondence: Professor Richard D. F. Harris, Xfi Centre for Finance and Investment, University of Exeter, Exeter EX4 4PU, UK. Tel: +44 (0) 1392 263215. Fax: +44 (0) 1392 263242. Email: R.D.F.Harris@exeter.ac.uk. Parts of this paper were written while Stoja was a visiting researcher at Georgetown University. He gratefully acknowledges financial support from Queen's University and would like to thank Martin Evans and Tim Bollerslev for useful comments and suggestions.

1. Introduction

There is now a large literature on the performance of conditional minimum variance hedging models. This literature is motivated by the now well-established fact that the conditional variance-covariance matrix of short horizon asset returns is both time-varying and highly predictable (see, for example, Andersen, Bollerslev and Diebold, 2004). When hedging a spot position with a position in the futures market, the minimum-variance hedge ratio (MVHR) is equal to the ratio of the covariance of spot and futures returns to the variance of futures returns. Exploiting the predictability in the conditional variance-covariance matrix of spot and futures returns should, in principle, lead to an improvement in hedging performance. In particular, the conditional MVHR minimizes the conditional variance of the hedge portfolio in each period, while the unconditional MVHR will, in any particular period, be either too low or too high (see, for example, Kroner and Sultan, 1993).

A common approach to conditional minimum-variance hedging is to model the time-varying conditional variance-covariance matrix of returns using, for example, a multivariate GARCH model, and use forecasts from this model to construct a forecast of the conditional MVHR. While the motivation for conditional minimum-variance hedging is clear, its benefits in practice appear to be limited. In particular, empirical evidence suggests that in terms of the reduction in variance of the hedged portfolio, conditional minimum-variance hedging at best offers only very marginal improvements over unconditional minimum-variance hedging. For example, Kroner and Sultan (1993) find that the variance of the conditionally hedged portfolio is between 2.2 percent and 4.6 percent lower than the variance of the unconditionally hedged portfolio. Brooks and Chong (2001) find that the variance of the conditionally hedged portfolio varies from 0.3 percent higher to 2.8 percent lower than the variance of the unconditionally hedged portfolio. After allowing for transaction costs that arise from the implementation of a time-varying hedging strategy, exploiting the time-variation in the conditional variance-covariance matrix of asset returns would in many cases lead to a *reduction* in the hedger's utility.[1]

[1] Of course, the actual impact on utility depends on (a) the size of the transaction costs and (b) the hedger's degree of risk aversion.

On the face of it, the results from the conditional minimum-variance hedging literature are difficult to reconcile with the results from the conditional volatility literature. On the one hand, conditional volatility models have been shown to explain as much as 50 percent of the variation in the integrated variance-covariance matrix of returns (see, for example, Andersen, Bollerslev, Diebold and Labys, 2003), while on the other hand, conditional minimum-variance hedging models – which use the same conditional volatility models in their construction – offer virtually no improvement over unconditional minimum-variance hedging models that counterfactually assume that the variance-covariance matrix of asset returns is constant. The interpretation of these findings, however, is far from straightforward. Indeed, there are several reasons why a conditional hedging model may fail to provide a significant improvement over the unconditional MVHR model. First, it may be that the unobserved 'integrated' MVHR is not time-varying to the extent that it is economically useful. In particular, it may be that while a significant benefit is obtained from estimating the unconditional or 'average' MVHR, there is only minimal advantage from forecasting the period-by-period deviations of the MVHR from this average. Second, it may be that while the integrated MVHR is time-varying, it is not predictable. In other words, the integrated MVHR is time-varying but the conditional MVHR is not. Third, it may be that the integrated MVHR is time-varying and predictable, but the conditional hedging model does not contain significant information about the integrated MVHR. In other words, it is simply ineffective as a forecasting model. Fourth, it may be that the conditional hedging model does contain significant information about the integrated MVHR but, owing to its misspecification, it incorporates this information inefficiently, in which case the MVHR forecasts that it generates will not yield the maximal reduction in hedged portfolio variance. In other words, for a conditional hedging model to be effective in terms of hedged portfolio variance reduction, it must be both informative *and* efficient. These are not merely questions of academic interest. For example, if we know that conditional hedging using an informative and efficient model cannot, even in principle, lead to significant reductions in portfolio variance then we can eschew conditional MVHR models in favour of the unconditional MVHR model. Conversely, if it is found that the poor performance of conditional hedging models stems from their systematic inefficiency, then we can focus our efforts on finding conditional MVHR models that are more efficient.

In this paper, we investigate the performance of conditional hedging models in the context of cross-hedging three daily exchange rates, USD/EUR, USD/GBP and USD/JPY, over the five year period 03/01/01 to 29/12/06. Our aim is not to compare the relative performance of individual conditional MVHR models (which has been done extensively elsewhere), but to characterize and explain their performance. In so doing, we employ the realized beta framework of Andersen, Bollerslev, Diebold and Wu (2005). Specifically, we construct an estimate of the integrated MVHR for each of the three exchange rate pairs, defined as the ratio of the integrated covariance of the hedged exchange rate and the hedging exchange rate to the integrated variance of the hedging exchange rate. The integrated MVHR represents the *ex post* upper bound in terms of hedging performance and therefore reveals the limit of the economic usefulness of conditional minimum-variance hedging.[2] It also provides a natural benchmark against which to evaluate the forecasting performance of conditional hedging models and to explain their underperformance. The integrated MVHR is estimated using the realized daily MVHR constructed from round-the-clock 30-minute returns.

Using the realized MVHR as a benchmark, we evaluate the out-of-sample performance of a number of MVHR models. The first is the unconditional MVHR, estimated as the OLS slope coefficient in a regression of the hedged currency return on the hedging currency return. This represents the lower bound for the conditional MVHR models. The second model is the *RiskMetrics* EWMA model in which the elements of the variance-covariance matrix are modelled as an exponentially weighted moving average of past squares and cross-products of returns. The third model is the diagonal vech multivariate GARCH model of Bollerslev, Engle and Wooldridge (1988). The fourth model is the constant correlation multivariate GARCH model of Bollerslev (1990), which assumes that the correlation between spot and futures returns is time-invariant. The fifth model is the S-GARCH model of Harris, Stoja and Tucker (2007), which imputes the conditional covariance of returns from the conditional variances of the sum of returns and the difference of returns. The sixth model is an ARMA (p, q) model in the realized MVHR. We first evaluate the performance of the

[2] By *ex post*, we mean conditioning on the contemporaneous integrated variance-covariance matrix, not conditioning on contemporaneous returns, which would clearly give rise to a perfect hedge.

conditional MVHR models in terms of hedging effectiveness, relative to the unconditional MVHR model. Consistent with the extant literature, we find that the conditional MVHR models essentially provide no improvement over the unconditional MVHR model. For example, for the USD/GBP hedged with the USD/JPY, the reduction in hedged portfolio variance using the unconditional MVHR is 27.7 percent. Using the conditional MVHR models, the reduction varies from 24.8 percent (for the *RiskMetrics* EWMA model) to 28.3 percent (for the ARMA model).

We next evaluate the performance of the conditional MVHR relative to the *ex post* upper bound. Using the realized MVHR, the average reduction in hedged portfolio variance is 36.9 percent, suggesting that conditional hedging may potentially be economically useful, but also that conditional MVHR models fall far short of achieving the *ex post* maximal reduction in hedged portfolio variance. In order to explain the poor performance of the conditional MVHR models, we evaluate their statistical performance from a forecasting perspective in terms of accuracy, bias and efficiency. We find that while conditional MVHR models provide virtually no improvement in terms of hedging performance, they do contain significant information about the true, time-varying MVHR. For example, the R-squared in a Mincer-Zarnowitz (1969) regression of the realized MVHR on the forecast MVHR is as high as 21 percent. However, in most cases, the conditional MVHR models are both biased (the MVHR forecasts that they generate are not correct on average) and, in some cases, inefficient (the MVHR forecasts that they generate are typically too volatile). This suggests that one possible explanation for the poor hedging performance of conditional MVHR models is that they are systematically mis-specified. Indeed, for some models, a decomposition of the mean square error (MSE) of the conditional MVHR models suggests that in some cases, as much as 13 percent of the forecast error is systematic rather than random.

To explore whether this systematic bias is able to account for the poor hedging performance of the conditional MVHR models, we use the estimated parameters of the Mincer-Zarnowitz regressions to construct forecasts of the MVHR that contain exactly the same information as the original MVHR forecasts, but which are unbiased and efficient by construction, and evaluate the hedging performance of these 'bias-corrected' MVHR forecasts. Using the bias-corrected forecasts, the average reduction in hedged portfolio variance increases by only about one percent. Even when the conditioning information set is expanded to include additional variables, hedging performance is still only marginally improved, suggesting that the poor performance of the conditional MVHR models is not attributable to their systematic bias and inefficiency. We therefore conclude that while systematic bias is partly to blame, most of the poor performance of conditional MVHR models can be attributed to unpredictability in the unobserved integrated MVHR.

The outline of the remainder of this paper is as follows. In Section 2 we discuss the theoretical framework for volatility and hedging. Section 3 describes the data, models and evaluation methodology. Section 4 presents the empirical results. Section 5 concludes.

2. Theoretical Framework

Consider two assets with continuous logarithmic prices, $p_1(t)$ and $p_2(t)$, whose sample paths constitute semi-martingales described by the stochastic differential equation

$$dp_i(t) = \mu_i(t)dt + \sqrt{\sigma_{ii}(t)}dW_i(t) \qquad i = 1,2 \qquad (1)$$

where $\mu_i(t)$ is the instantaneous drift, $\sigma_{ii}(t)$ is the instantaneous variance, and $W_i(t)$ are correlated geometric Brownian motion processes with instantaneous correlation $\rho_{12}(t)$, instantaneous covariance $\sigma_{12}(t) = \rho_{12}(t)\sqrt{\sigma_{11}(t)}\sqrt{\sigma_{22}(t)}$ and $\text{cov}(dW_i(t), \sigma_{ij}(t)) = 0$ for $i, j = 1,2$. Suppose that prices are observed at discrete intervals $t = 1,\ldots,T$. The stochastic process governing the discretely observed logarithmic returns, defined as $r_{i,t} = p_i(t) - p_i(t-1)$, is given by

$$r_{i,t} = \mu_{i,t} + z_t\sqrt{\sigma_{ii,t}}, \qquad i = 1,2 \qquad (2)$$

where z_t is a zero-mean, unit-variance white noise process. The elements of the integrated variance-covariance matrix of $r_{1,t}$ and $r_{2,t}$ are given by

$$\sigma_{ij,t} = \int_{t-1}^{t} \sigma_{ij}(s)ds, \qquad i, j = 1,2 \tag{3}$$

(See, for example, Andersen and Bollerslev, 1998). Consider now an investor who chooses to hedge a long position in Asset 1 with a short position in the correlated Asset 2 over one period. The time-varying integrated MVHR can be defined as

$$h_{12,t} = \frac{\sigma_{12,t}}{\sigma_{2,t}^{2}} \tag{4}$$

The integrated MVHR, $h_{12,t}$, is not observable but a consistent estimate of it is given by

$$h_{12,t}^{RV(q)} = \frac{\hat{\sigma}_{12,t}^{RV(q)}}{\hat{\sigma}_{22,t}^{RV(q)}} \tag{5}$$

where $\sigma_{ij,t}^{RV(q)} = \sum_{s=1}^{1/q} r_{i,t-1+sq}^{q} r_{j,t-1+sq}^{q}$ is the realized covariance between asset i and asset j and $1/q$ is an integer that represents the sampling frequency of intraday returns, $r_{i,t}^{q} = p_i(t) - p_i(t-q)$. Recent theoretical research has shown that under very general conditions $\sigma_{ij,t}^{RV(q)}$ converges uniformly in probability to $\sigma_{ij,t}$ as $q \to 0$ (see, for example, the extended discussion in Andersen, Bollerslev and Diebold, 2004). Since the elements of the realized variance-covariance matrix used in (5) are continuous-record consistent estimators of the corresponding elements of the integrated variance-covariance matrix given by (3), the realized MVHR, $h_{12,t}^{RV(q)}$ is also a continuous-record consistent estimator of the underlying integrated MVHR, $h_{12,t}$ (see Andersen, Bollerslev, Diebold and Wu, 2005, 2006, who consider the analogous problem of estimating the realized CAPM beta for individual stocks).

In practice, the accuracy of the realized volatility approach is limited by the fact that market microstructure effects distort the measurement of returns at high frequencies in such a way that measured returns no longer satisfy the regularity conditions that are required for the consistency properties of realized volatility. In particular, microstructure effects induce an upward bias in estimated volatility that increases as the measurement interval becomes smaller (see, for example, Ait-Sahalia, Mykland and Zhang, 2005; Zhang, Ait-Sahalia and Mykland, 2003; Bandi and Russell, 2003). Consequently, some researchers have proposed estimation of integrated volatility by sampling returns at non-negligible time intervals. Empirical evidence suggests that intervals between five and 30 minutes are effective for the estimation of integrated volatility (Andersen, Bollerslev, Diebold, and Labys, 2001, 2003; Barndorff-Nielsen and Shephard, 2002, 2004).

In implementing minimum-variance hedging in practice, we require a *forecast* of the integrated MVHR, or, equivalently, an estimate of the *conditional* MVHR, given by

$$\hat{h}_{12,t} = E[h_{12,t} \mid \Lambda_{t-1}] \qquad (6)$$

where $E[.]$ is the conditional expectation and Λ_t is the time-t information set. Many conditional MVHR models have been proposed, most of which are based on models of conditional volatility, such as multivariate EWMA or multivariate GARCH (see, for example, Lien and Tse, 2002). The efficacy of a conditional MVHR model depends on a number of factors. The first is the degree of predictability in the integrated MVHR (or equivalently, the degree of time-variation in the conditional MVHR). From the extensive literature on volatility modeling, it is clear that the elements of the integrated variance-covariance matrix are both time-varying and highly predictable. However, this does not necessarily imply that the integrated MVHR is also highly predictable. In particular, Andersen, Bollerslev, Diebold and Wu (2005) note that the elements of the integrated variance-covariance matrix potentially share common features, and so the dynamic properties of non-linear functions of those elements may be very different from the dynamic properties of the elements themselves. For example, in their study of the integrated CAPM beta, Andersen, Bollerslev, Diebold and Wu (2005) find that while integrated variances and covariances are highly persistent, and best modelled as a fractionally integrated process, the integrated CAPM beta is much less persistent, and best modelled as a stationary process. Clearly, in the context of hedging, this has important implications for the predictability of the integrated MVHR, and the limits to the performance of conditional MVHR models.

The second factor that determines the effectiveness of a conditional MVHR model is the validity of the model itself. The conditional volatility models on which conditional MVHR models are usually based are unlikely to accurately capture the true data generating process of returns, and so they can, at best, be assumed to be *ad hoc* approximations that are essentially mis-specified (see, for example, Nelson and Foster, 1994). Thus, any model of the conditional MVHR is also likely to be mis-specified in the sense that it will incorporate information about the integrated MVHR in an inefficient way. Furthermore, being a non-linear function of conditional variances and covariances, the conditional MVHR may potentially be more mis-specified than the conditional volatility model on which it is based.[3] Mis-specification will result in ineffective estimates of the conditional MVHR, even though the model itself may contain substantial information about the unobserved integrated MVHR. In other words, for a conditional MVHR model to be effective, it must be both informative *and* efficient.

The third factor that determines the effectiveness of a conditional MVHR model is the economic usefulness of conditional minimum-variance hedging. It may simply be that the ability to forecast period-by-period deviations of the time-varying MVHR from the average MVHR is of limited incremental value. In other words, hedging with the integrated MVHR is not significantly more effective than hedging with the unconditional MVHR. In this case, a conditional MVHR model might be informative and efficient with respect to the integrated MVHR, but will not generate a substantial reduction in hedged portfolio variance. In this paper, we investigate these issues.

[3] In particular, any bias that may be present in the constituents of the MVHR may be magnified through Jensen's inequality. For a discussion of biases in non-linear transformations of the conditional variance-covariance matrix, see Harris and Guermat (2006).

3. Data, Models and Evaluation Criteria

We characterize and explain the hedging performance of a number of conditional MVHR models in cross-hedging three currencies, each measured against the USD. These currencies are USD/GBP, USD/EUR and USD/JPY. Specifically, we assume that a USD investor has a long position in a foreign currency and hedges the risk exposure of this position using a short position in one of the remaining two currencies. While in many cases it would clearly be more effective to hedge a long currency position using currency futures, there are situations where currency cross-hedging may be appropriate. For example, many corporations have exposure to two or more currencies simultaneously. An efficient approach to hedging this exposure is to first exploit the natural cross-hedge that arises from the non-zero correlation between the different currency exposures, and then to use derivatives to hedge the residual risk.[4] Moreover, cross-hedging is relevant to the activities of a currency hedge fund that is, for example, long in an undervalued currency and short in an overvalued currency. In any case, however, the cross-hedging example that we use serves as a useful illustration for other hedging situations, irrespective of the exact nature of the hedged portfolio. We assume that the investor undertakes 'inventory' hedging by minimising the variance of daily hedged portfolio returns. In this section, we describe in detail the conditional MVHR models, the data used in their estimation and the criteria by which they are evaluated.

3.1 Data

We use intraday data for each of the three exchange rates over the period 03 January 2001 to 29 December 2006. The data, which were provided by Bank of America, are 30-minute exchange rates for the three currencies against the USD. The market operates around the clock and so there are a total of 48 observations each day, or 75,072 observations in total for each series. The data contained three outliers that were clearly the result of data entry errors and so the values for these observations

[4] This is borne out by survey evidence. See "Survey of Derivatives Usage by US Non-financial Firms", conducted by the Weiss Center for International Financial Research, Wharton School, and CIBC World Markets, 1998.

were linearly interpolated from the adjacent exchange rates. For the benchmark in the empirical analysis, the raw data were used to calculate 30-minute log returns and these were used to construct the realized MVHR for each exchange rate pair (*i*, *j*) as

$$h_{ij,t}^{RV(1/48)} = \frac{\sigma_{ij,t}^{RV(1/48)}}{\sigma_{jj,t}^{RV(1/48)}}, \qquad i, j = 1,2 \tag{7}$$

where

$$\sigma_{ij,t}^{RV(1/48)} = \sum_{s=1}^{48} r_{i,t-1+s/48}^{(1/48)} r_{j,t-1+s/48}^{(1/48)}, \qquad i, j = 1,2 \tag{8}$$

Rather than directly using the realized MVHR, we apply a linear correction that maximizes the in-sample performance of the realized MVHR. In particular, we construct the corrected realized MVHR

$$\tilde{h}_{ij,t}^{RV(1/48)} = a_{ij} + b_{ij} h_{ij,t}^{RV(1/48)} \tag{9}$$

For each pair of currencies, we choose a_{ij} and b_{ij} to minimize the variance of the hedged portfolio over the forecast period. Such a correction helps to mitigate the measurement error in the raw realized MVHR.[5] Note that the linearly corrected realized MVHR contains exactly the same information as the raw realized MVHR but, by construction, yields the lowest possible hedged portfolio variance and hence represents the appropriate benchmark against which to evaluate the conditional

[5] The difference between the performance of the raw realized MVHR and the corrected realized MVHR is minimal, and the choice between the two does not influence the conclusions that we draw. We also experimented with various other corrections designed to reduce the impact of measurement error. One was the use of the Kalman filter to extract the integrated MVHR from a state space equation with the measurement equation specified with the asymptotically correct measurement error calculated from the intraday data (see Andersen, Bollerslev, Diebold and Wu, 2005). However, the filtered MVHR was found to be substantially *worse* than the realized MVHR in terms of hedging performance. Another approach was to construct the realized MVHR including one lead and one lag in the covariance equation, helping to reduce any measurement error that is due to non-synchronous trading. Again, however, this approach worsened the performance of the realized MVHR.

MVHR models. The 30-minute returns were summed from 00.30am to 12.00 midnight to yield daily log returns for use in estimating the conditional volatility models. The full sample is divided into an initial estimation period from 02 January 2001 to 02 December 2002 (500 observations), and a forecast period from 03 December 2002 to 29 December 2006 (1,064 observations). Table 1 reports summary statistics for the daily log returns and the realized variances, covariances and MVHRs over the forecast period.

[Table 1]

Summary statistics for the three daily log return series are reported in Panel A. For all three currencies, daily log returns are negatively skewed and leptokurtic, although for the GBP and EUR, the departure from normality is not that severe. The LM(4) test for up to fourth order serial correlation suggests that the return series are approximately serially uncorrelated. The ARCH(4) test for up to fourth order serial correlation in squared returns shows that all three currency return series display significant volatility clustering, motivating the use of conditional MVHR models. The three return series are moderately correlated, with the highest correlation between GBP and EUR, and the lowest between GBP and JPY. The realized variances for the three currencies and the realized covariances between them are reported in Panel B. As expected, all six series are highly time-varying and non-normal (see, for example, Andersen, Bollerslev, Diebold and Labys, 2001). The LM(4) test suggests that there is a high degree of linear predictability in the realized variances and covariances, which is consistent with the ARCH(4) test for the return series. Summary statistics for the three realized MVHR series are reported in Panel C. Like the elements of the realized variance-covariance matrix, the realized MVHRs for the three currencies are highly time-varying and display significant serial correlation. However, note that the degree of time-variation in the realized MVHR is somewhat lower than for the realized variances and covariance. For example, the coefficient of variation (the standard deviation divided by the mean) is between 0.65 and 1.08 for the realized variances and covariances, but between 0.32 and 0.57 for the realized MVHR. Nevertheless, the existence of predictable time-variation in the realized MVHR implies that conditional minimum-variance hedging may potentially be economically useful.

3.2 Hedging Models

The lower benchmark for the conditional MVHR models is the unconditional MVHR, measured as the OLS slope coefficient in a regression of the hedged exchange rate on the hedging exchange rate, estimated over the forecast sample. The upper benchmark is represented by the realized MVHR given by (9). We consider five conditional MVHR models. Four of these (the EWMA, diagonal vech, constant correlation and S-GARCH models) are based on models of the conditional variance-covariance matrix. The fifth model is a dynamic model of the realized MVHR. Each of the conditional MVHR models is re-estimated each day to provide a one-step ahead forecast of the integrated MVHR. The conditional MVHR models that we employ are as follows:

RiskMetrics EWMA Model

The *RiskMetrics* EWMA model of JP Morgan (1994) specifies the conditional variance or covariance to be exponentially weighted moving averages of the squares and cross-products of returns, respectively. In its recursive form, the EWMA model is given by

$$\hat{\sigma}_{ij,t} = \lambda \hat{\sigma}_{ij,t-1} + (1-\lambda) r_{i,t-1} r_{j,t-1}, \qquad i, j = 1,2 \tag{10}$$

where λ is the decay factor. The model is implemented using the *RiskMetrics* decay factor of 0.94. With a single decay factor, the EWMA model generates a conditional variance-covariance matrix that is positive semi-definite by construction.

The Diagonal Vech GARCH Model

The diagonal vech multivariate GARCH(1, 1) model of Bollerslev, Engle and Wooldridge (1988) is a multivariate generalisation of the univariate GARCH(1, 1) model of Bollerslev (1986) and is given by

$$r_{i,t} = \mu_i + \varepsilon_{i,t}, \qquad i = 1,2 \tag{11}$$

$$\hat{\sigma}_{ij,t} = \alpha_{ij,0} + \alpha_{ij,1}\hat{\sigma}_{ij,t-1} + \alpha_{ij,2}\varepsilon_{i,t-1}\varepsilon_{j,t-1}, \qquad i, j = 1,2 \qquad (12)$$

The model is estimated as a bivariate system for each pair of currencies. Unlike the EWMA model, positive semi-definiteness of the conditional variance-covariance matrix is not guaranteed. However, for all three currency pairs, the estimated correlation coefficients were found to be between –1 and +1 for all observations.

The Constant Correlation GARCH Model

The constant correlation multivariate GARCH model of Bollerslev (1990) assumes that the covariance between two asset returns is proportional to the product of their respective standard deviations. The model is given by

$$r_{i,t} = \mu_i + \varepsilon_{i,t}, \qquad i = 1,2 \qquad (13)$$

$$\hat{\sigma}_{ij,t} = \begin{cases} \alpha_{ii,0} + \alpha_{ii,1}\hat{\sigma}_{ii,t-1} + \alpha_{ii,2}\varepsilon^2_{i,t-1} & \text{for } i = j \\ \hat{\rho}_{ij}(\hat{\sigma}_{ii,t}\hat{\sigma}_{jj,t})^{0.5} & \text{for } i \neq j \end{cases}, \qquad i, j = 1,2 \qquad (14)$$

where $\hat{\rho}_{ij}$ is the time-invariant correlation coefficient. The conditions required to ensure a positive semi-definite conditional variance-covariance matrix are easily imposed, since it only requires that $-1 \leq \hat{\rho}_{ij} \leq 1$. The constant correlation model is estimated as a bivariate system for each pair of currencies.

The Simplified Multivariate GARCH Model

The S-GARCH of Harris, Stoja and Tucker (2007) involves the estimation of only univariate GARCH models, both for the individual return series and for the sum and difference of each pair of series. The covariance between each pair of return series is then imputed from these conditional variance estimates. First, the conditional variances are estimated using univariate GARCH(1, 1) models:

$$r_{i,t} = \mu_i + \varepsilon_{i,t}, \qquad i = 1,2 \qquad (15)$$

$$\hat{\sigma}_{ii,t} = \alpha_{ii,0} + \alpha_{ii,1}\hat{\sigma}_{ii,t-1} + \alpha_{ii,2}\varepsilon^2_{i,t-1}, \qquad i = 1,2 \tag{16}$$

Then the new auxiliary variables, $r_{+,t} = r_{1,t} + r_{2,t}$ and $r_{-,t} = r_{1,t} - r_{2,t}$, are constructed, and univariate GARCH(1, 1) models used to estimate the conditional variances of these.

$$r_{i,t} = \mu_i + \varepsilon_{i,t}, \qquad i = +,- \tag{17}$$

$$\hat{\sigma}_{ii,t} = \alpha_{ii,0} + \alpha_{ii,1}\hat{\sigma}_{ii,t-1} + \alpha_{ii,2}\varepsilon^2_{i,t-1}, \qquad i = +,- \tag{18}$$

The conditional covariance between each pair of currencies is then imputed using the identity

$$\sigma_{ij,t} \equiv (1/4)(\sigma^2_{+,t} - \sigma^2_{-,t}) \tag{19}$$

Like the diagonal vech model, the S-GARCH model does not guarantee that the conditional variance-covariance matrix is positive semi-definite. However, for all three currency pairs, the estimated correlation coefficients were again found to be between –1 and +1 for all observations.

The ARMA(p, q) *Model*

Finally, the realized MVHR is modelled as an ARMA(p, q) process given by

$$h^{RV(1/48)}_{ij,t} = \alpha_{ij,0} + \sum_{k=1}^{p}\alpha_{ij,k}h^{RV(1/48)}_{ij,t-k} + \sum_{l=1}^{q}\beta_{ij,l}\varepsilon_{ij,t-l} + \varepsilon_{ij,t} \tag{20}$$

The initial estimation sample was used to determine the optimal lag structure, which led to the choice of an ARMA(1, 1) model for all three currency pairs. This specification was maintained for the out-of-sample forecast period, but the model parameters were re-estimated each day.

Estimation issues

The diagonal vech, constant correlation and S-GARCH models were estimated by Quasi Maximum Likelihood function with a Gaussian conditional distribution, using the BHHH algorithm. For the diagonal vech model, there were a small number of cases (less than one percent) where the estimation did not converge. For these cases, the model parameters were set to the values estimated in the previously converged iteration. For the constant correlation and S-GARCH models, every iteration converged. The ARMA model was estimated using the Gauss-Newton algorithm with numerical derivatives. For the EWMA model, the initial values for each element of the conditional variance-covariance matrix are set equal to the corresponding elements of the unconditional variance-covariance matrix, estimated over the initial estimation sample.

3.3 Evaluation

To evaluate the performance of the different models across the three cross-hedged currency portfolios, we consider both economic and statistical criteria. We first examine the hedging effectiveness of the conditional MVHR models. Specifically, we use the estimated conditional MVHR for each pair of currencies to construct a hedged portfolio whose return is given by

$$r_{ij,t} = r_{i,t} - \hat{h}_{ij,t} r_{j,t} \tag{21}$$

We calculate the percentage reduction in the variance of the hedged portfolio relative to the variance of the unhedged currency. Our measure of hedging effectiveness is therefore given by

$$HE(\hat{h}_{ij,t}) = \frac{\text{var}(r_{ij,t}) - \text{var}(r_{i.t})}{\text{var}(r_{i,t})} \tag{22}$$

We next analyze the statistical performance of the conditional MVHR models from a forecasting perspective. Our primary measure of statistical performance is the mean square error of the estimated conditional MVHR with respect to the linearly corrected realized MVHR, given by

$$MSE(\hat{h}_{ij,t}) = \frac{1}{T}\sum_{t=1}^{T}(\tilde{h}_{ij,t}^{RV(1/48)} - \hat{h}_{ij,t})^2 \qquad (23)$$

To investigate the statistical properties of the conditional MVHR models, we employ a Mincer-Zarnowitz (1969) regression for each model given by

$$\tilde{h}_{ij,t}^{RV(1/48)} = \alpha_{ij} + \beta_{ij}\hat{h}_{ij,t} + \varepsilon_{ij,t} \qquad (24)$$

where $E(\varepsilon_{ij,t}) = 0$ and $E(\varepsilon_{ij,t}\hat{h}_{ij,t}) = 0$. Since we are interested in the *ex post* performance of the (*ex ante*) conditional MVHR estimates, the regression is estimated in-sample over the forecast period to yield a single set of parameter estimates for each conditional MVHR model, for each currency pair. The estimated parameters from the Mincer-Zarnowitz regression measure the degree of bias and inefficiency in the conditional MVHR estimates. In particular, a necessary and sufficient condition for $\hat{h}_{ij,t}$ to be unconditionally unbiased is that $a_{ij} = (1 - b_{ij})E(\tilde{h}_{ij,t}^{RV(1/48)})$ (see, for example, Taylor, 1999, who considers this regression in the context of volatility forecasting). A necessary condition for $\hat{h}_t$ be weakly efficient is that $a_{ij} = 0$ and $b_{ij} = 1$. For the unbiasedness hypothesis, we use the sample average of $\tilde{h}_{ij,t}^{RV(1/48)}$, evaluated over the forecast sample, as an unbiased and consistent estimate of its expected value. The *R*-squared from the Mincer-Zarnowitz regression measures the information content of the estimated conditional MVHR, irrespective of any bias or inefficiency. The Mincer-Zarnowitz regression allows us to explore the potential reasons for the poor performance of the conditional MVHR models. In particular, we can use the results of the Mincer-Zarnowitz regression to decompose the MSE of a conditional MVHR model into its systematic and non-systematic components in the following way:

$$
\begin{aligned}
MSE(\hat{h}_{ij,t}) &= \frac{1}{T}\sum_{t=1}^{T}(\tilde{h}_{ij,t}^{RV(1/48)} - \hat{h}_{ij,t})^2 \\
&= (\overline{\tilde{h}}_{ij}^{RV(1/48)} - \overline{\hat{h}}_{ij})^2 + (1-b_{ij})^2 \operatorname{var}(\hat{h}_{ij,t}) \\
&\quad + (1-R_{ij}^2)\operatorname{var}(\tilde{h}_{ij,t}^{RV(1/48)})
\end{aligned} \tag{25}
$$

where $\overline{\hat{h}}_{ij} = (1/T)\sum_{t=1}^{T}\hat{h}_{ij,t}$ and $\overline{\tilde{h}}_{ij}^{RV(1/48)} = (1/T)\sum_{t=1}^{T}\tilde{h}_{ij,t}^{RV(1/48)}$. The first two terms represent the proportion of the MSE that can be explained by the bias and inefficiency of the conditional MVHR model, respectively, and together, these represent the systematic component of the MSE. The last term represents the component of the MSE that is non-systematic, or random.

Having established the statistical properties of the conditional MVHR models, we evaluate the extent to which these properties impact their hedging performance. In particular, we use the estimated parameters of the Mincer-Zarnowitz regression to construct 'bias-adjusted' estimates of the conditional MVHR as

$$
\hat{\hat{h}}_{ij,t} = \hat{a}_{ij} + \hat{b}_{ij}\hat{h}_{ij,t} \tag{26}
$$

We could expect the bias-corrected estimate, $\hat{\hat{h}}_{ij,t}$, to perform better empirically than the uncorrected estimate $\hat{h}_{ij,t}$, since it contains the same information as $\hat{h}_t$ (in the sense that the *R*-squared in the Mincer-Zarnowitz regression will be the same whether $\hat{\hat{h}}_{ij,t}$ or $\hat{h}_{ij,t}$ is used as a regressor), but will, by construction, be both unbiased and weakly efficient. We investigate the performance of the bias-corrected MVHR forecasts using the hedging effectiveness measure given by (22).

The preceding Mincer-Zarnowitz regression is a test of *weak* form efficiency since it conditions only on the forecast MVHR itself. We also test a stronger form of efficiency by estimating an augmented Mincer-Zarnowitz regression that includes additional conditioning variables:

$$\tilde{h}_{ij,t}^{RV(1/48)} = \alpha_{ij} + b_{ij}\hat{h}_{ij,t} + c_{ij}'X_{t-1} + \varepsilon_{ij,t} \tag{27}$$

where $E(\varepsilon_{ij,t}) = 0$, $E(\varepsilon_{ij,t}\hat{h}_{ij,t}) = 0$ and X_{t-1} is a vector of variables that are in the time $t - 1$ information set, Λ_{t-1}. A necessary condition for $\hat{h}_t$ be strongly efficient is that $a_{ij} = 0$, $b_{ij} = 1$ and $c_{ij} = 0$. In X_{t-1}, we include $\tilde{h}_{ij,t-1}^{RV(1/48)}$, together with its square and cross-product with $\hat{h}_{ij,t}$, up to the third lag. We adopt a general-to-specific testing methodology to eliminate insignificant variables to yield a final model that contains only significant variables. As with the previous case, we use the estimated parameters of the augmented Mincer-Zarnowitz regression to construct bias-corrected forecasts, given by

$$\hat{\hat{h}}_{ij,t}^{Aug} = \hat{a}_{ij} + \hat{b}_{ij}\hat{h}_{ij,t} + \hat{c}_{ij}'X_{t-1} \tag{28}$$

We assess the performance of $\hat{\hat{h}}_{ij,t}^{Aug}$ using the hedging effectiveness measure given by (22).

4. Results

Table 2 reports the percentage reduction in variance relative to the unhedged currency for the three hedged portfolios, for the unconditional MVHR model and the five conditional MVHR models. On average, the conditional MVHR models perform very poorly. The average reduction across the five models is 60.3 percent, 27.3 percent and 29.3 percent, respectively, for GBP-EUR, GBP-JPY for EUR-JPY. This compares with 60.2 percent, 27.7 percent and 30.5 percent for the unconditional model. For the GBP-JPY and EUR-JPY portfolios, the ARMA model provides the best hedging performance of 28.3 percent and 30.3 percent, respectively, although the latter is still lower than the reduction provided by the unconditional model. For the GBP-EUR portfolio, the constant correlation model provides the greatest reduction of 60.4 percent, which is marginally better than the reduction provided by the unconditional model. There is very little difference between the performance of the three GARCH models, although the constant correlation model appears to be marginally better than the diagonal vech and S-GARCH models. The EWMA model is consistently the worst, with an average reduction in hedge portfolio variance that is considerably lower than the reduction provided by the unconditional model. Consistent with the extant literature, therefore, we find that conditional MVHR models are not able to provide significant improvements in terms of hedging performance over the unconditional MVHR model. Indeed, if transaction costs were accounted for, these results would almost certainly imply a reduction in hedger utility even for the best performing models.

The last line of Table 2 gives the reduction in hedged portfolio performance for the realized MVHR. To the extent that the realized MVHR is an accurate estimate of the unobserved integrated MVHR, this represents the *ex post* maximal reduction in hedged portfolio variance. Not only do conditional MVHR models fail to improve upon the unconditional MVHR model, they also fall far short of achieving this maximal reduction. For example, for EUR-JPY, the average reduction in variance across the five conditional MVHR models is 29.3 percent, while for the unconditional MVHR model, it is 30.5 percent. Using the realized MVHR, the reduction in hedged portfolio variance is 38.3 percent. Thus, we can conclude that the integrated MVHR is not only time-varying, but in a way that is economically significant. However, conditional MVHR models do not appear to capture the time-variation in the integrated MVHR in a way that is economically useful.

[Table 2]

Table 3 reports the results of the Mincer-Zarnowitz regression given by (24). The table reports the estimated slope and intercept, together with standard errors in parentheses, and the regression R-squared. It also reports the F-statistic to test the null hypothesis of weak efficiency, i.e. that the intercept is zero and the slope is one. In all cases, we can reject the null hypothesis of efficiency very strongly. For GBP-JPY and EUR-JPY, the rejection is strongest for the EWMA model, and weakest for the ARMA model, while for GBP-EUR, the converse is true. In almost all cases, the intercept is significantly different from zero, implying that the conditional MVHR forecasts are incorrect on average. For the GBP-JPY and EUR-JPY portfolios, the slope coefficient is significantly different from one and in most cases, less than one, implying that the conditional MVHR forecasts are too dispersed. The bias and inefficiency is particularly bad for the EUR-JPY portfolio, where in many cases, the slope parameter is not significantly different from zero.

While the forecasts of the conditional MVHR models are generally biased and inefficient, it is clear from the Mincer-Zarnowitz regressions that they nevertheless, in some cases at least, contain significant information about the unobserved integrated MVHR. In particular, the average R-squared across the five models is 0.19, 0.16 and 0.03, for the GBP-EUR, GBP-JPY and EUR-JPY portfolios, respectively. Indeed, the EWMA model, which has the greatest degree of bias and inefficiency, also has the highest average R-squared across the three currencies. Even for the EUR-JPY portfolio, the EWMA model explains more then 10 percent of the variation in the realized MVHR. In contrast, the ARMA model explains only five percent of the variation for this portfolio, and the three GARCH models explain none of it. The highest R-squared overall is 21 percent, provided by the ARMA model for the GBP-EUR portfolio. Thus, it is clear that in spite of the bias and inefficiency inherent in all of the conditional MVHR models, in many cases they nevertheless contain significant information about the unobserved integrated MVHR.

[Table 3]

To shed further light on the degree of systematic bias in the conditional MVHR models, Table 4 reports the results of the decomposition of the MSE into the components due to bias, inefficiency and noise, for each of the five models, for each of the three currency pairs. The systematic component accounts for between just under two percent of MSE (for the ARMA model for EUR-JPY) to over 13 percent of MSE (for the EWMA model for EUR-JPY). On average across the three currency pairs, the systematic component of MSE is largest for the EWMA model and smallest for the ARMA model.

[Table 4]

Thus, one possible explanation for the poor performance of conditional MVHR models is that they are systematically biased and inefficient. To evaluate the impact that the systematic bias reported in Table 3 has on the hedging performance of the five conditional MVHR models, we use the estimated parameters of the Mincer-Zarnowitz regression reported in Table 3 to construct bias-corrected MVHR forecasts given by (26), and use these to construct a hedge portfolio for each currency pair. The bias-corrected MVHR forecasts contain the same information as the uncorrected MVHR forecasts but are, by construction, unbiased and weakly efficient. Panel A of Table 5 reports the reduction in variance for the bias-corrected MVHR forecasts. In all cases, the bias-correction improves the hedging performance of the conditional MVHR models, and in some cases, the improvements are substantial. For example, for the EWMA model, the reduction in hedge portfolio variance improves from 24.8 percent to 27.0 percent for GBP-JPY and from 26.7 percent to 30.2 percent for EUR-JPY. As expected, the greatest improvement is seen for the conditional MVHR models for which the systematic bias was the greatest component of MSE (the EWMA model) and the smallest improvement for the model for which the systematic component was smallest (the ARMA model). On average across the five conditional MVHR models, the bias correction increases the reduction in hedge portfolio variance from 60.3 percent to 61.0 percent for GBP-EUR, from 27.3 percent to 28.1 percent for GBP-JPY, and from 29.3 percent to 30.3 percent for EUR-JPY. However, even after the bias-adjustment, this represents only a marginal improvement over the variance reduction provided by the unconditional MVHR model (60.2 percent, 27.7 percent and 30.5 percent, respectively), and still falls far short of the variance reduction obtained using the realized MVHR (65.2 percent, 36.9 percent and 38.3 percent, respectively). Thus, while the systematic bias generally accounts for a considerable fraction of the MSE for the conditional MVHR models, correcting for this systematic bias does not significantly improve their hedging performance relative to the unconditional MVHR model.

[Table 5]

The preceding analysis suggests that the poor performance of conditional hedging models is not due to weak inefficiency. A test of strong form efficiency would condition on the entire information set, not just the conditional MVHR. While such a test of strong form efficiency is not possible, we can include additional variables in the conditioning information set. In particular, we estimate the ‘augmented’ Mincer-

Zarnowitz regression given by (27) that includes $\hat{h}_{ij,t}$ and $\tilde{h}_{ij,t-1}^{RV(1/48)}$, together with the squares and cross-products of these up to the third lag. We adopt a general-to-specific testing methodology to eliminate insignificant variables to yield a final model that contains only significant variables. We use the parameters of the augmented Mincer-Zarnowitz regression to construct the bias-corrected forecasts given by (28), and use these to hedge the currency portfolio. Panel B of Table 5 reports the percentage reduction in variance relative to the unhedged currency for each model and each currency portfolio using these augmented bias-corrected MVHR forecasts.[6] As expected, the inclusion of additional conditioning variables in the Mincer-Zarnowitz regression improves the performance of the bias-adjusted MVHR forecasts. However, the improvement is small, and the conditional MVHR models still fall far short of achieving the maximal reduction in hedge portfolio variance represented by the realized MVHR.

5. Summary and Conclusion

It is now widely accepted that the conditional variance-covariance matrix of short horizon asset returns is both highly time-varying and highly predictable. This has led to the development of conditional MVHR models that are widely used both in academia and in practice. However, empirical evidence suggests that these conditional MVHR models offer only very marginal improvements in terms of hedge ratio performance relative to the unconditional MVHR model that assumes, counterfactually, that the variance-covariance matrix of returns is constant. Indeed, after allowing for transaction costs, conditional minimum-variance hedging would in many cases lead to a reduction in hedger utility. In this paper, we explore the possible reasons for this finding using the realized beta framework of Andersen, Bollerslev, Diebold and Wu (2005).

We first show that the hedging performance of conditional MVHR models, in addition to providing only marginal improvement over the unconditional MVHR model, fall far short of the *ex post* maximal reduction in hedge portfolio variance that

[6] For brevity, we do not report the estimation results for the augmented Mincer-Zarnowitz regression, nor the corresponding MSE decomposition. These results are available from the authors on request.

is obtained using the realized MVHR. We find that conditional MVHR models are generally systematically biased (the forecasts that they generate are on average incorrect) and, in many cases, weakly inefficient (the forecasts that they generate are too dispersed), and that this systematic bias and inefficiency accounts for as much as 13 percent of the forecast error relative to the realized MVHR. Correcting for this bias and inefficiency improves the performance of conditional MVHR forecasts, but it still falls far short of that obtained by using the realized MVHR. Inclusion of additional conditioning variables in the bias-correction offers only a marginal further improvement. We therefore conclude that the poor performance of conditional MVHR models relative to the unconditional MVHR model, although partly accounted for by systematic bias and inefficiency, is largely due to unpredictability in the unobserved integrated MVHR. A natural explanation for the fact that the integrated MVHR is less predictable than the elements of the integrated variance-covariance matrix is that there are common features in the integrated variance-covariance matrix. These common features mean that certain non-linear functions of volatility are less persistent, and hence less predictable, than the volatility measures on which they are based.

The results of this analysis have important implications for the users of hedging models. In particular, they suggest that in spite of the extensive efforts that have been made in the academic literature to find effective models of the time-varying conditional MVHR, the models proposed do not provide a significant improvement over the unconditional MVHR, which counterfactually assumes that the variance-covariance matrix of returns is constant. Moreover, while the poor performance of conditional MVHR models can be partly attributed to the fact that they are mis-specified (lending hope to the possibility that better specified models can be found), most of the variation in the integrated MVHR is simply unpredictable.

References

Aït-Sahalia, Y., P. A. Mykland, and L. Zhang, 2005, "How often to sample a continuous-time process in the presence of market microstructure noise", *Review of Financial Studies,* Vol 18, 2, 351-416.

Andersen, T.G., and T. Bollerslev, 1998, "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts", *International Economic Review* 39, 885-905.

Andersen, T.G., T. Bollerslev and F.X. Diebold, 2004, "Parametric and nonparametric measurements of volatility". In: Aït-Sahalia, Y., Hansen, L.P. (Eds.), *Handbook of Financial Econometrics*, North-Holland, Amsterdam.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2001, "The Distribution of Exchange Rate Realized Volatility", *Journal of the American Statistical Association* 96, 42–55.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2003, "Modeling and forecasting realized volatility", *Econometrica* 71, 579–625.

Andersen, T.G., T. Bollerslev, F. Diebold and J. Wu, 2006, "A Framework for Exploring the Macroeconomic Determinants of Systematic Risk", *American Economic Review* 95, 398-404.

Andersen, T.G., T. Bollerslev, F. Diebold and J. Wu, 2006, "Realized Beta: Persistence and Predictability" in T. Fomby and D. Terrell (eds.) *Advances in Econometrics: Econometric Analysis of Economic and Financial Time Series in Honor of R.F. Engle and C.W.J. Granger* , Volume B, 1-40.

Bandi, F., and J. Russell. 2003, "Microstructure Noise, Realized Volatility, and Optimal Sampling", Working Paper. University of Chicago. Graduate School of Business.

Barndorff-Nielsen, O.E., and N. Shephard, 2002, "Estimating quadratic variation using realized variance", *Journal of Applied Econometrics* 17, 457–477.

Barndorff-Nielsen, O.E., and N. Shephard, 2004a, "Econometric analysis of realized covariation: High frequency based covariance, regression and correlation in financial economics", *Econometrica* 72, 885–925.

Barndorff-Nielsen, O.E.. and Shephard, N. 2004b, "A Feasible Central Limit Theory for Realized Volatility Under Leverage", *Manuscript*, Nuffield College, Oxford.

Bollerslev, T., 1986, "Generalized Autoregressive Conditional Heteroscedasticity", *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.

Brooks, C., and J. Chong, 2001, “The Cross-Currency Hedging Performance of Implied versus Statistical Volatility Models”, *Journal of Futures Markets*, Vol. 21(11), 1043–1069.

Harris, R.D.F., and C. Guermat, (2006), “Bias in the Estimation of Non-Linear Transformations of the Latent Variance of Returns”, *Journal of Forecasting*, 25, 481-494.

Harris, R.D.F., E. Stoja and J. Tucker, 2007, “A Simplified Approach to Modeling the Comovement of Asset Returns”, *Journal of Futures Markets*, Vol. 27(6), 575–598.

JP Morgan, 1996, *RiskMetrics*™ Technical Document, 4th Edition, New York.

Kroner, K., and J. Sultan, 1993, “Time Varying Distribution and Dynamic Hedging with Foreign Currency Futures”, *Journal of Finance and Quantitative Analysis* 28, 535-551.

Lien, D., and Y.K. Tse, 2002, “Some Recent Developments in Futures Hedging”, *Journal of Economic Surveys*, Vol. 16(3), 357-96.

Mincer, J., and V. Zarnowitz, 1969, “The Valuation of Economic Forecasts,” in *Economic Forecasts and Expectations*, Ed. J. Mincer, National Bureau of Economic Research.

Nelson, D., and D. Foster, 1994, “Asymptotic Filtering Theory for Univariate ARCH Models”, *Econometrica* 62, 1-41.

Taylor, J., 1999, “Evaluating Volatility and Interval Forecasts”, *Journal of Forecasting* 18, 111-128.

Zhang, L., Y. Aït-Sahalia and P.A. Mykland, 2005, “A tale of two time scales: Determining integrated volatility with noisy high-frequency data”, *Journal of the American Statistical Association*, 100 1394-1411.

Table 1 Summary Statistics

Panel A: Returns

	GBP	EUR	JPY
Mean (%)	0.02	0.03	0.00
Stand Dev (%)	0.54	0.60	0.55
Skewness	-0.054	0.009	0.146
Kurtosis	0.455	0.552	1.422
B-J	9.664	13.508	93.389
LM(4)	1.578	6.109	4.294
ARCH(4)	648.580	657.330	723.180

Correlations

	GBP	EUR	JPY
GBP	1.000	0.776	0.527
EUR		1.000	0.552
JPY			1.000

Notes: The table reports the mean standard deviation, skewness and kurtosis for daily log close-to-close returns (Panel A), realized variances and covariances (Panel B) and the realized MVHR (Panel C), for USD/GBP, USD/EUR and USD/JPY for the forecast period 03/12/02 to 29/12/06 (1064 daily observations). The five percent critical values of the B-J, LM(4) and ARCH(4) statistics are 5.99, 9.49 and 9.49, respectively.

Table 1 Summary Statistics (Continued)

Panel B: Realized Variances and Covariances

	GBP	**EUR**	**JPY**
Mean	2.80E-05	3.42E-05	3.29E-05
Stand Dev	1.81E-05	2.38E-05	2.82E-05
Skewness	2.588	2.650	4.951
Kurtosis	10.643	12.784	38.587
B-J	6203.246	8483.231	70293.299
LM(4)	243.775	206.981	101.812
	GBP-EUR	**GBP-JPY**	**EUR-JPY**
Mean	2.17E-05	1.30E-05	1.68E-05
Stand Dev	2.34E-01	1.32E-05	1.48E-05
Skewness	1.700	3.654	3.544
Kurtosis	23.978	31.940	27.191
B-J	25978.352	47548.706	34972.353
LM(4)	159.525	105.160	83.300

Notes: The table reports the mean standard deviation, skewness and kurtosis for daily log close-to-close returns (Panel A), realized variances and covariances (Panel B) and the realized MVHR (Panel C), for USD/GBP, USD/EUR and USD/JPY for the forecast period 03/12/02 to 29/12/06 (1064 daily observations). The five percent critical values of the B-J and LM(4) statistics are 5.99 and 9.49, respectively.

Table 1 Summary Statistics (Continued)

Panel C: Minimum-Variance Hedge Ratio

	GBP	EUR	JPY
Mean	0.697	0.557	0.649
Stand Dev	0.234	0.319	0.323
Skewness	1.700	-0.011	0.095
Kurtosis	23.978	1.350	1.071
B-J	25978.352	80.708	52.407
LM(4)	357.643	289.176	322.509

Notes: The table reports the mean standard deviation, skewness and kurtosis for daily log close-to-close returns (Panel A), realized variances and covariances (Panel B) and the realized MVHR (Panel C), for USD/GBP, USD/EUR and USD/JPY for the forecast period 03/12/02 to 29/12/06 (1064 daily observations). The five percent critical values of the B-J and LM(4) statistics are 5.99 and 9.49, respectively.

Table 2 Percentage Reduction in Hedged Portfolio Variance (in percent)

	GBP-EUR	GBP-JPY	EUR-JPY
Unconditional	60.16	27.74	30.49
EWMA	60.20	24.80	26.74
DVech	60.20	27.94	29.90
CCor	60.40	27.80	29.95
S-GARCH	60.31	27.85	29.74
ARMA	60.31	28.30	30.25
Realized	65.18	36.93	38.25

Notes: The table reports the percentage reduction in the variance of the hedged portfolio return over the forecast period, relative to the unhedged currency.

Table 3 Mincer-Zarnowitz Regression Results

	GBP-EUR				GBP-JPY				EUR-JPY			
	α	β	R^2	F	α	β	R^2	F	α	β	R^2	F
EWMA	0.149	0.769	0.200	22.131	-0.206	0.664	0.182	35.378	-0.352	0.475	0.106	80.097
	(0.032)	(0.045)			(0.025)	(0.043)			(0.028)	(0.042)		
DVech	0.085	0.967	0.178	54.220	-0.082	1.063	0.160	75.802	-0.599	0.092	0.000	64.478
	(0.038)	(0.060)			(0.034)	(0.075)			(0.082)	(0.148)		
CCor	0.028	1.036	0.191	34.695	-0.074	1.080	0.160	74.989	-0.668	-0.033	0.000	51.794
	(0.041)	(0.062)			(0.035)	(0.076)			(0.094)	(0.165)		
S-GARCH	0.027	1.042	0.191	39.261	-0.084	1.063	0.155	77.746	-0.717	-0.119	0.001	59.257
	(0.040)	(0.062)			(0.035)	(0.076)			(0.089)	(0.155)		
ARMA	0.058	0.993	0.206	35.046	-0.069	0.993	0.164	27.214	-0.211	0.711	0.052	10.500
	(0.039)	(0.060)			(0.035)	(0.069)			(0.058)	(0.093)		

Notes: The table reports the estimated intercept, slope and R-squared of the Mincer-Zarnowitz regression given by equation (24) in the main text. Standard errors for the estimated parameters are reported in parentheses. The table also reports the F-statistic for the null hypothesis that the intercept is equal to zero and the slope is equal to one. The five percent critical value of the F-statistic is 3.00.

Table 4 Mean Square Error Decomposition

	GBP-EUR				GBP-JPY				EUR-JPY			
Unconditional	0.054	-	-	-	0.103	-	-	-	0.106	-	-	-
EWMA	0.045	0.56%	2.19%	97.34%	0.089	0.93%	5.33%	93.84%	0.107	0.51%	12.62%	86.96%
DVech	0.049	8.36%	0.02%	91.70%	0.098	12.44%	0.06%	87.58%	0.117	7.66%	3.18%	89.25%
CCor	0.047	5.63%	0.03%	94.44%	0.098	12.29%	0.09%	87.70%	0.114	5.52%	3.38%	91.19%
S-GARCH	0.047	2.74%	0.04%	93.74%	0.099	12.73%	0.06%	87.30%	0.116	5.64%	4.41%	90.04%
ARMA	0.046	6.20%	0.00%	93.89%	0.089	4.88%	0.00%	95.21%	0.101	1.05%	0.89%	98.15%

Notes: The table reports the mean square error, given by (23), and the decomposition of the mean square error, given by (25), into the systematic components that are due to bias and inefficiency, and the random component.

Table 5 Percentage Reduction in Hedged Portfolio Variance (in percent)

Panel A: Corrected for Bias and Weak Inefficiency

	GBP-EUR	**GBP-JPY**	**EUR-JPY**
Unconditional	-60.16	-27.74	-30.49
EWMA	-60.66	-26.99	-30.17
DVech	-61.09	-28.54	-30.32
CCor	-61.01	-28.33	-30.32
S-GARCH	-60.98	-28.44	-30.35
ARMA	-60.94	-28.40	-30.38
Realized	-65.18	-36.93	-38.25

Notes: The table reports the percentage reduction in the variance of the hedged portfolio return over the forecast period, relative to the unhedged currency, using the bias-corrected forecasts given by (26).

Table 5 Percentage Reduction in Hedged Portfolio Variance (in percent) (Continued)

Panel B: Corrected for Bias and Strong Inefficiency

	GBP-EUR	GBP-JPY	EUR-JPY
Unconditional	-60.16	-27.74	-30.49
EWMA	61.05	27.54	30.48
DVech	61.29	28.85	30.82
CCor	61.15	28.52	30.83
S-GARCH	61.09	28.61	30.75
ARMA	61.30	28.30	30.87
Realized	-65.18	-36.93	-38.25

Notes: The table reports the percentage reduction in the variance of the hedged portfolio return over the forecast period, relative to the unhedged currency, using the bias-corrected forecasts given by (28).