NCER Working Paper Series

Volatility timing and portfolio selection: How best to forecast volatility

Adam Clements
Annastiina Silvennoinen

Working Paper #76
October 2011

Volatility timing and portfolio selection: How best to forecast volatility

A Clements and A Silvennoinen
School of Economics and Finance, Queensland University of Technology, NCER.

Abstract

Within the context of volatility timing and portfolio selection this paper considers how best to estimate a volatility model. Two issues are dealt with, namely the frequency of data used to construct volatility estimates, and the loss function used to estimate the parameters of a volatility model. We find support for the use of intraday data for estimating volatility which is consistent with earlier research. We also find that the choice of loss function is important and show that a simple mean squared error loss, overall provides the best forecasts of volatility upon which to form optimal portfolios.

Keywords
Volatility, volatility timing, utility, portfolio allocation, realized volatility

JEL Classification Numbers
C22, G11, G17

Corresponding author
Adam Clements
Professor in Finance School of Economics and Finance
Queensland University of Technology
GPO Box 2434, Brisbane, 4001
Qld, Australia

email `a.clements@qut.edu.au`
+61 7 3138 2525
Annastiina Silvennoinen
Post-Doctoral Research Fellow
School of Economics and Finance
Queensland University of Technology
GPO Box 2434, Brisbane, 4001
Qld, Australia

email `silvennoinen@qut.edu.au`
+61 7 3138 2920

1 Introduction

Forecasts of volatility are important inputs into numerous financial applications, including derivative pricing, risk estimation and portfolio allocation. The modern volatility forecasting literature stems from the seminal work of Engle (1982) and Bollerslev Bollerslev (1986) in a univariate setting, and from Bollerslev Bollerslev (1990) and Engle (2002) among others in the multivariate setting. For a broad overview of the major developments in this field, see Campbell, Lo, and Mackinlay (1997), Gourieroux and Jasiak (2001) and Andersen, Bollerslev, Christoffersen, and Diebold (2006).

A voluminous literature exists dealing with modeling of, and forecasting volatility. Much of this literature examines the relative performance of competing forecasts in a generic statistical setting, that is without any consideration of an economic application of the forecasts. For a wide ranging overview of such literature see (Poon and Granger, 2003, 2005), or for a more comprehensive comparison of forecasts, see Hansen and Lunde (2005) or Becker and Clements (2008). Relatively speaking, there is less literature that considers the economic value of forecasting volatility (volatility timing) within the context of portfolio allocation. Graham and Harvey (1996) and Copeland and Copeland (1999) study trading rules based on changes in volatility. West, Edison, and Cho (1993) undertake a utility based comparison of the economic value of a range of volatility forecasts. Fleming, Kirby, and Ostdiek. (2001) examine the value of volatility timing in the context of a short horizon asset allocation strategy. To do so, they consider a mean-variance investor allocating wealth across stocks, bonds and gold based on forecasts of the variance-covariance matrix of returns.

In recent years there have been many developments in the measurement of volatility by utilizing high frequency intraday data, a principle stemming from the earlier work of Schwert (1989). Andersen, Bollerslev, Diebold, and Labys (2001), Andersen, Bollerslev, Diebold, and Labys (2003) and Barndorff-Nielsen and Shephard (2002) among others advocate the use of realized volatility as a more precise estimate of volatility relative to those based on lower frequency data[1]. Fleming et al. (2003) build upon Fleming et al. (2001) and highlight the positive economic value of realized volatility relative to estimates of volatility based on daily returns.

Traditionally, volatility models such as GARCH models, along with those based on realized volatility are estimated by quasi-maximum likelihood (QML). Parameter estimates obtained under a QML loss function are used to subsequently generate forecasts applied in the portfolio allocation context, see (Fleming et al., 2001, 2003). Skouras (2007) proposes a different approach

[1]Following Fleming, Kirby, and Ostdiek. (2003) we use the general realized volatility term to refer to the full realized covariance matrix of asset returns. In later sections, we refer specifically to variances, covariances and correlation

in which a utility based metric is used to estimate the parameters of a univariate volatility model. Such an approach has much to recommend it as the criteria under which the model is estimated and then applied are consistent.

While there is no doubt that realized volatility offers a superior estimate of volatility, we do not understand whether this superiority is influenced by the loss function under which a volatility model is estimated. In other words, is realized volatility still preferred if alternatives to QML are used. This paper considers how to best estimate a volatility model that will be used for the purposes of portfolio allocation. The problem will be considered along two dimensions, volatility estimates based on daily or intraday data, and the loss function under which the model is estimated. Along with QML, both utility based and simple mean squared error loss functions are examined. A three asset, portfolio allocation problem involving equities (S&P 500: SP), Treasury notes (10 year Treasury bonds: TY) and gold (GC) will be examined. To take advantage of recent econometric advances, the model chosen is the MIDAS approach of Ghysels, Santa-Clara, and Valkanov (2005).

We confirm the findings of Fleming et al. (2003) in that realized volatility is of positive economic value. We also find that the choice of loss function is important. When using realized volatility, simple mean squared error and QML are equivalent with utility being inferior. It is also found mean squared error produces the most stable forecasts and hence portfolio exposures, an important issue when considering transactions costs.

The paper proceeds as follows. Section 2 outlines the general portfolio allocation framework along with how model performance will be compared. Section 3 outlines the volatility model considered and the competing loss functions under which estimation occurs. Section 4 describes the data employed, with Section 5 outlining the empirical results. Section 6 provides concluding comments.

2 The portfolio allocation problem and forecast evaluation

We follow Skouras (2007) and consider an investor with negative exponential utility,

$$u(r_{p,t}) = -\exp(-\lambda\, r_{p,t}) \tag{1}$$

where r_p is the portfolio return realized by the investor during the period to time t and λ is their coefficient of risk aversion.

We assume the vector of excess returns $\boldsymbol{r}_t$ obey

$$\boldsymbol{r}_t \sim \boldsymbol{F}(\boldsymbol{\mu}, \boldsymbol{\Sigma}_t)\,, \tag{2}$$

where $\boldsymbol{F}$ is some multivariate distribution, $\boldsymbol{\mu}$ is fixed vector of expected excess returns and $\boldsymbol{\Sigma}_t$ is the conditional covariance matrix of returns. The manner in which the portfolio of risky assets will be constructed is now described.

Begin by defining $\overline{\boldsymbol{\Sigma}}_t$ as a forecast of the conditional covariance matrix, $\mathbf{w_t}$ as a vector of portfolio weights and μ_0 to be the target return for the portfolio. The composition of the optimal portfolio is then given by

$$\mathbf{w_t} = \frac{\overline{\boldsymbol{\Sigma}}_t^{-1}\boldsymbol{\mu}}{\boldsymbol{\mu}'\overline{\boldsymbol{\Sigma}}_t^{-1}\boldsymbol{\mu}}\,\mu_0\,. \tag{3}$$

Portfolio returns are then determined by $r_{p,t} = \mathbf{w_t}'\boldsymbol{r}_t$. The manner in which $\overline{\boldsymbol{\Sigma}}_t$ is obtained is described in the following section.

We follow (Fleming et al., 2001, 2003) in comparing the performance of the various forecasts in terms of the relative economic benefit they produce when used to form optimal portfolios. We find a constant, Δ that solves

$$\sum_{t=1}^{T} U(r_{p,t}^1) = \sum_{t=1}^{T} U(r_{p,t}^2 - \Delta) \tag{4}$$

where $r_{p,t}^1$ and $r_{p,t}^2$ represent portfolio returns based on two competing estimation criteria. Here Δ reflects the incremental value of using the second approach as opposed to the first. It measures the maximum return an investor would be willing to sacrifice, on average per day, to capture the gains of switching to the second criteria.

The portfolio choice implied by equation 3 requires estimates of the vector of expected returns, $\boldsymbol{\mu}$. To control for the uncertainty surrounding the expected returns, the block bootstrap approach of Fleming et al. (2003) is used. Artificial samples of 10000 observations are generated by randomly selecting blocks of random length, with replacement from the original sample[2]. Mean returns for the three assets are computed from the artificial sample and are used as an estimate for $\boldsymbol{\mu}$ in equation 3. We use the same constraints on the expected returns as Fleming et al. (2003). A bootstrap is acceptable if $\mu_{SP} > \mu_{TY} > \mu_{GC}$, $\mu_{SP} > 0$, $\mu_{TY} > 0$ and $\sigma_{SP} > \sigma_{TY}$, where μ and σ denote the sample mean and standard deviation of the artificial returns. Given a series of volatility forecasts, $\overline{\boldsymbol{\Sigma}}_t\,\forall t = 1, \ldots, T$, for each acceptable bootstrap, portfolio weights (and hence portfolio returns) are computed from equation 3 using the expected returns obtained from that specific bootstrap. From the utilities generated by two volatility forecasts, the difference in economic value Δ is computed for each bootstrap. This procedure is repeated 500 times with a mean value for Δ across the 500 bootstraps reported in annualized basis points below.

[2]Smaller sample sizes of 2000 and 5000 were also used. There is no qualitative difference to results presented here.

3 Forecasting the covariance matrix

The approach for generating a forecast of the covariance matrix, $\overline{\mathbf{\Sigma}}_t$ is drawn from the recent advances in MIDAS regressions. This methodology produces volatility forecasts directly from a weighted average of past observations of volatility. Following from Ghysels et al. (2005) a forecast of the conditional covariance matrix, $\overline{\mathbf{\Sigma}}_t$ is generated by

$$\overline{\mathbf{\Sigma}}_t = \sum_{k=1}^{k_{\max}} b\left(k, \boldsymbol{\theta}\right) \widehat{\mathbf{\Sigma}}_{t-k} \tag{5}$$

where $\widehat{\mathbf{\Sigma}}_{t-k}$ are historical estimates of the covariance matrix (the estimates used here will be described below). In this instance, the same scalar MIDAS weights, $b\left(k, \boldsymbol{\theta}\right)$ will be applied to all elements of $\widehat{\mathbf{\Sigma}}_{t-k}$ for each lag k. The maximum lag length $k_{\max}$ can be chosen rather liberally as the weight parameters $b\left(k, \boldsymbol{\theta}\right)$ are tightly parameterized. All subsequent analysis is based on $k_{\max} = 100$. Here the weights are determined by means of a beta density function and normalized such that $\sum b\left(k, \boldsymbol{\theta}\right) = 1$. A beta distribution function is fully specified by the 2×1 parameter vector $\boldsymbol{\theta}$. Here $\theta_1 = 1$ meaning that only the θ_2 must be estimated. The constraint $0 < \theta_2 < 1$ ensures that the weighting function is a decreasing function of the lag k. The historical estimates of the covariance matrix will now be described

$\widehat{\mathbf{\Sigma}}_t$: Daily returns (D)

The estimate of the covariance matrix is simply the outer-product of the vector of daily returns, $\boldsymbol{r}_t$

$$\widehat{\mathbf{\Sigma}}_t = \boldsymbol{r}_t \boldsymbol{r}_t^{'}. \tag{6}$$

$\widehat{\mathbf{\Sigma}}_t$: Intraday returns (RV)

Here the estimate is the sum of the outer-product of intraday returns,

$$\widehat{\mathbf{\Sigma}}_t = \sum_{i=1}^{N} \boldsymbol{r}_t^i \boldsymbol{r}_t^{i'} \tag{7}$$

where N represents the number of intraday intervals.

The loss functions under which values for θ_2 in equation (5) is estimated will be described.

Quasi- Maximum Likelihood QML

The value for θ_2 is chosen so as to

$$\underset{\theta_2}{\operatorname{argmax}} \sum_{t=1}^{T} log(|\overline{\mathbf{\Sigma}}_t|) + \boldsymbol{r}_t \overline{\mathbf{\Sigma}}_t^{-1} \boldsymbol{r}_t'. \tag{8}$$

Minimum Mean Squared Error MSE

Under this estimation criteria, θ_2 is chosen so as to

$$\underset{\theta_2}{\operatorname{argmin}} \sum_{t=1}^{T} \operatorname{vec}(\overline{\boldsymbol{\Sigma}}_t - \widehat{\boldsymbol{\Sigma}}_t)' \operatorname{vec}(\overline{\boldsymbol{\Sigma}}_t - \widehat{\boldsymbol{\Sigma}}_t). \tag{9}$$

Utility Based Estimation UTL

Skouras (2007) proposes a method by which the parameters of a univariate volatility model can be estimated directly within an economic criteria. As opposed to likelihood maximization, Skouras (2007) suggests estimating parameters by maximizing the utility realized from the portfolios formed from model forecasts.

Given the optimal portfolio rule in equation (3), and the expression for realized utility in equation (1), the objective function for a maximum utility estimator is

$$\underset{\theta_2}{\operatorname{argmax}} \frac{1}{T} \sum_{t=1}^{T} - \exp(-\lambda \mathbf{w_t}' \boldsymbol{r}_t). \tag{10}$$

Parameter estimation is conducted on the basis of optimally weighting historical volatility so as to construct portfolios that lead to the greatest expected utility as opposed to statistically optimal forecasts of volatility.

4 Data

The portfolio allocation problem considered here relates to a mix of bond, equities and gold. The study treats returns on S&P 500 Composite Index futures as equities exposure (SP), returns on U.S. 10-year Treasury Note futures as bond market exposure (TY) along with returns on Gold futures (GC) [3]. Data was gathered for the period covering 1 July 1997 to 29 June 2009 giving a sample of 2985 observations. RV estimates of the covariance matrix were constructed by summing the cross products of 15 minute futures contract returns.

Table 1 reports the means and standard deviations of each of the individual elements in the covariance matrix. Panel A reports on the covariance estimates where $\widehat{\boldsymbol{\Sigma}}_t$: D and Panel B, $\widehat{\boldsymbol{\Sigma}}_t$: RV. It is clear from comparing the means, that on average both the daily (D) and RV estimates are virtually indistinguishable from each other. A comparison of the standard deviations in the right hand columns reflect the well know pattern that more precise volatility estimates are obtained by using high frequency intraday returns. In most cases the standard deviations of the daily based volatility estimates are 50% larger than the corresponding RV estimates. The greater precision of RV normally leads to improved forecasts and portfolio outcomes. This

[3]Intraday data for both futures contracts were purchased from Tick Data

	Mean$\times 10^{-3}$				Standard Deviation$\times 10^{-3}$		
			Panel A: $\widehat{\boldsymbol{\Sigma}}_t$: D				
	SP	TY	GC		SP	TY	GC
SP	0.1741	-0.0153	-0.0037	SP	0.6223	0.0944	0.2428
TY	-	0.0171	0.0042	TY	-	0.0389	0.0644
GC	-	-	0.1328	GC	-	-	0.3551
			Panel B: $\widehat{\boldsymbol{\Sigma}}_t$: RV				
	SP	TY	GC		SP	TY	GC
SP	0.1712	-0.0170	-0.0020	SP	0.4455	0.0666	0.1591
TY	-	0.0159	0.0028	TY	-	-	0.0397
GC	-	-	0.1366	GC	-	-	0.3452

Table 1: Means and standard deviations of the individual elements in the historical estimates of the covariance matrix. Panel A show results for $\widehat{\boldsymbol{\Sigma}}_t$ given daily returns, and Panel B intraday returns.

widely held belief will be reconsidered in light of altering the loss function under which volatility forecasts are generated.

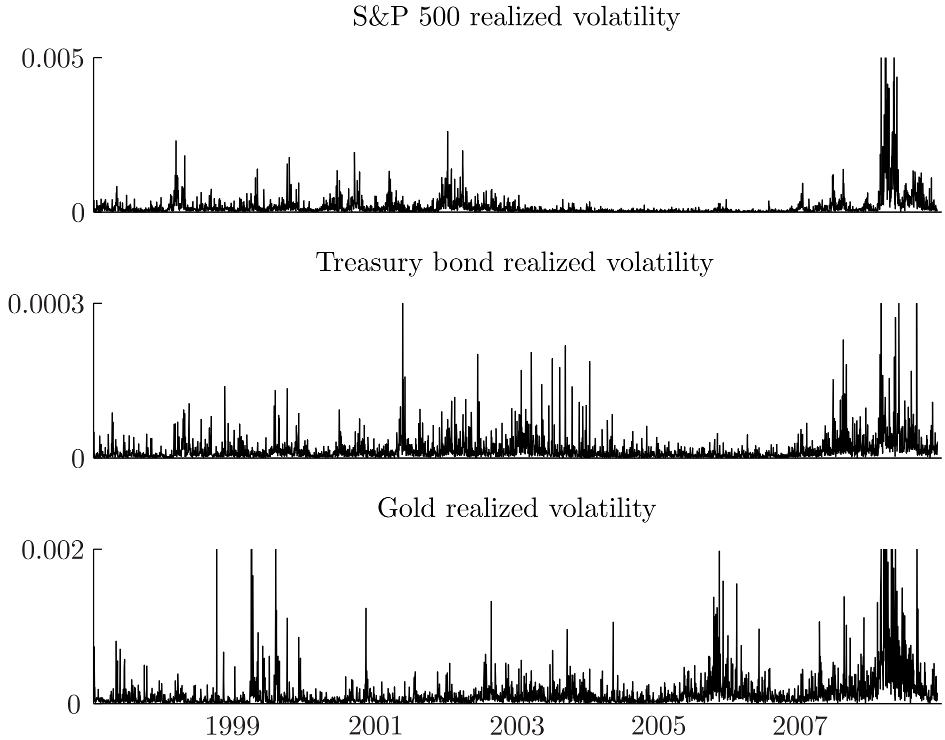

Figure 1: S&P 500 RV estimates (top panel), Treasury bond RV estimates (middle panel) and Gold RV estimates (bottom panel).

Figures 1 and 2 plot the realized volatilities and correlations of the three assets considered[4]. Figure 1 shows the realized volatility of equity futures (top panel), bond futures returns (middle

[4]Correlations are shown as they are easier to interpret. Volatility estimates using daily returns are not shown as they are simply a more noisy estimate relative to the RVs shown here.

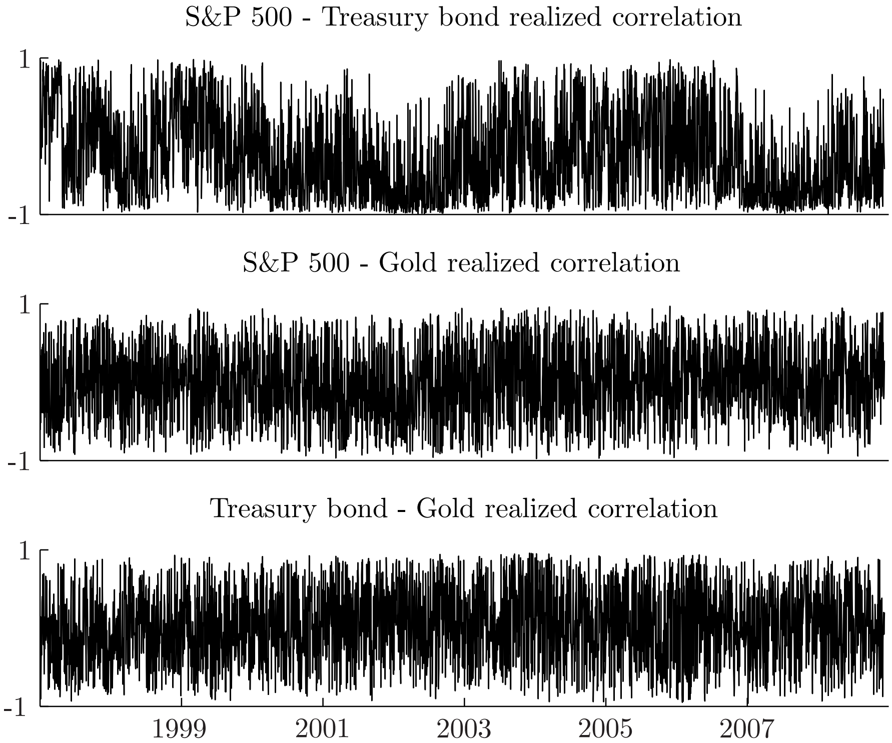

Figure 2: S&P 500 and RV Treasury bond realized correlation estimates (top panel), S&P 500 and Gold realized correlation estimates (middle panel) and Treasury bond and Gold realized correlation estimates (bottom panel).

panel) and gold futures returns (lower panel). Equity volatility shows a familiar pattern, low volatility during much of the sample period with higher volatility due to collapse of technology stocks. It is clear that the events surrounding the credit crisis of the second half of 2008 dominate in terms of the levels of volatility reached (the scale of the plot has been constrained otherwise no variation is evident due to the level of recent volatility). The volatility of bond returns is unsurprisingly much lower in magnitude than equity returns and generally more stable. It is evident that the recent financial crisis has lead to a sustained period of somewhat higher volatility. Volatility in gold returns rose in late 2005 and early 2006 due to central bank activity, and rose to historically high levels due the height of the recent market turmoil.

Realized correlations between the respective pairs of assets are shown in Figure 2. The correlation between equities and bonds (top panel) is quite persistent over time. It shows a downward trend through to 2002-2003 with it subsequently being weak during 2004-2006, followed by a period very strong negative correlation during much of the recent crisis. In contrast to the bond and equity case, neither the correlation between either equities and gold (middle panel) nor bonds and gold (lower panel) show any long-term persistence or structure.

5 Empirical results

Given the 2985 daily observations, the first 1000 observations were used as an initial estimation period. One day ahead forecasts of $\overline{\boldsymbol{\Sigma}}_t$ are obtained for $t = 1001$ and a portfolio formed according to Section 2 which leads to 1985 forecasts of $\overline{\boldsymbol{\Sigma}}_t$ and subsequent portfolio allocations. Here θ_2 is re-estimated every 200 days. Under the MSE and QML loss functions, θ_2 is not re-estimated for each bootstrap as $\boldsymbol{\mu}$ does not enter into the loss functions as they only require the covariance proxy (MSE) and daily returns (QML). However, in the case of UTL, θ_2 must be re-estimated as each new estimate of $\boldsymbol{\mu}$ obtained from each bootstrap enters directly into the loss function (equation 10) through the portfolio weights $\mathbf{w_t}$ obtained from equation 3.

Tables 2 through 5 report estimates of average Δ (averaged across 500 bootstraps) computed from equation 4. These represent the incremental economic value of using the forecast (a combination of loss function and $\widehat{\boldsymbol{\Sigma}}_t$, D or RV) in the column heading over that in row heading. These are expressed in annual basis point terms. The lower number in each cell corresponds to the proportion of bootstraps where Δ was found to be positive.

	MSE^D	MSE^{RV}	QML^D	QML^{RV}	UTL^D	UTL^{RV}
MSE^D	-	97.788 1.000	−46.616 0.000	95.267 1.000	−43.798 0.060	−37.992 0.092
MSE^{RV}		-	−153.707 0.000	−2.286 0.000	−150.388 0.000	−144.222 0.000
QML^D			-	136.373 1.000	1.874 0.094	7.239 0.632
QML^{RV}				-	−147.963 0.000	−141.845 0.000
UTL^D					-	4.872 0.606
UTL^{RV}						-

Table 2: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 6\%$, $\gamma = 2$. Each cell reports the average value for Δ across the 500 bootstraps and the proportion of bootstraps where Δ was found to be positive.

To begin, we will focus on Table 2 which reports results for the case of $\mu_0 = 6\%$ and $\gamma = 2$. In comparing MSE and QML, two patterns emerge. It is clear that RV is preferred to D when estimating $\widehat{\boldsymbol{\Sigma}}_t$ as MSE^{RV} is preferred MSE^D, along with QML^{RV} over QML^D. This is consistent with the finding of Fleming et al. (2003) who show that moving from daily, to RV estimates using intraday data leads to positive economic value. This is not surprising given the results in Table 1 in that the RV estimates are much less noisy than those based on daily returns. If one was restricted to using daily returns to compute $\widehat{\boldsymbol{\Sigma}}_t$, MSE is preferred over QML. On the other hand, if RV estimates are available (given the required intraday data) there is little difference between MSE and QML. While no value of Δ is positive out of the 500 bootstraps when QML^{RV} is compared to MSE^{RV}, the average value is only -2.286 basis points. Thus while MSE^{RV} is statistically inferior to QML^{RV}, from an economic point of view, there is little difference between the value of the two loss functions when RV is used. This finding has interesting implications for how a volatility model is estimated. From equation 8 is clear that the inverse of $\overline{\mathbf{\Sigma}}_t$ must be taken. Computationally this becomes a difficult issue when the dimension of the portfolio becomes large which in a practical setting is often the case. The fact that (when using RV) MSE leads to virtually the same economic outcome is an important finding in that larger dimensional problems can be handled using MSE avoiding the computational issues associated with QML. UTL^D and UTL^{RV} are inferior to both MSE^D and MSE^{RV} but approximately equivalent in performance to QML^D. Skouras (2007) finds that UTL is superior to QML, but only in a univariate setting using daily returns and does not consider the performance of MSE.

	MSE^D	MSE^{RV}	QML^D	QML^{RV}	UTL^D	UTL^{RV}
MSE^D	-	107.288 1.000	-64.115 0.000	103.514 1.000	-58.377 0.064	-54.370 0.104
MSE^{RV}		-	-194.298 0.000	-3.069 0.000	-186.364 0.000	-180.766 0.000
QML^D			-	154.743 1.000	1.567 0.156	2.635 0.624
QML^{RV}				-	-183.413 0.000	-178.382 0.000
UTL^D					-	0.729 0.658
UTL^{RV}						-

Table 3: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 6\%$, $\gamma = 5$. Each cell reports the average value for Δ across the 500 bootstraps and the proportion of bootstraps where Δ was found to be positive.

	MSE^D	MSE^{RV}	QML^D	QML^{RV}	UTL^D	UTL^{RV}
MSE^D	-	133.101 1.000	-67.618 0.000	129.350 1.000	-62.510 0.064	-55.221 0.084
MSE^{RV}		-	-217.189 0.000	-3.316 0.000	-211.023 0.000	-202.988 0.000
QML^D			-	187.330 1.000	3.115 0.110	9.407 0.634
QML^{RV}				-	-207.531 0.000	-199.634 0.000
UTL^D					-	5.565 0.596
UTL^{RV}						-

Table 4: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 8\%$, $\gamma = 2$. Each cell reports the average value for Δ across the 500 bootstraps and the proportion of bootstraps where Δ was found to be positive.

Tables 3 through 5 report average Δ for increasing μ_0 and, or γ. Overall, the results in terms of rankings discussed earlier remain the same. The differences in economic value reflected in the average value for Δ become greater in absolute terms, a pattern also observed in Fleming et al. (2003). There are three combinations where there is little difference between the economic value of the forecasts, QML^{RV} and MSE^{RV}, UTL^D and QML^D and UTL^D and UTL^{RV}. In these cases only relatively minor changes in average Δ occur as μ_0 or γ increase and hence the forecasts continue to perform in a very similar manner.

	MSE^D	MSE^{RV}	QML^D	QML^{RV}	UTL^D	UTL^{RV}
MSE^D	-	150.680 1.000	-97.206 0.000	144.810 1.000	-87.724 0.056	-86.935 0.112
MSE^{RV}		-	-288.366 0.000	-4.491 0.000	-274.193 0.000	-269.534 0.000
QML^D			-	219.788 1.000	0.030 0.200	-8.938 0.536
QML^{RV}				-	-270.470 0.000	-267.978 0.000
UTL^D					-	-8.565 0.660
UTL^{RV}						-

Table 5: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 8\%$, $\gamma = 5$. Each cell reports the average value for Δ across the 500 bootstraps and the proportion of bootstraps where Δ was found to be positive.

In practice, an investor faces transaction costs as they alter their portfolio holdings through time. These costs are a function of both the frequency and magnitude of portfolio changes. Here we do not take a stance on the form of the transaction costs but compare the mean absolute changes in portfolio weights and their standard deviations, to reveal whether a link exists between the competing forecasts and portfolio stability. These statistics are reported in Table 6. There is little difference between whether D or RV estimates are used under the MSE loss function. The magnitude and volatility of portfolio changes are quite similar for both MSE^D and MSE^{RV} across the three asset classes. However, when one moves to QML or UTL, the magnitude and volatility of changes in weights generally increase by 50% to nearly 100%. Overall, MSE (irrespective of the data used to estimate $\widehat{\boldsymbol{\Sigma}}_t$) produces the most stable forecasts and hence portfolio weights, which in practice minimizes transaction costs.

	MSE^D	MSE^{RV}	QML^D	QML^{RV}	UTL^D	UTL^{RV}
$\lvert\Delta_{w_{SP}}\rvert$	0.0303	0.0318	0.0570	0.0384	0.0524	0.0518
$\sigma_{\Delta_{w_{SP}}}$	0.0364	0.0381	0.0642	0.0454	0.0613	0.0602
$\lvert\Delta_{w_{TY}}\rvert$	0.0487	0.0514	0.0905	0.0619	0.0815	0.0827
$\sigma_{\Delta_{w_{TY}}}$	0.0568	0.0590	0.0988	0.0701	0.0931	0.0930
$\lvert\Delta_{w_{GC}}\rvert$	0.0344	0.0332	0.0633	0.0398	0.0570	0.0568
$\sigma_{\Delta_{w_{GC}}}$	0.0420	0.0384	0.0729	0.0452	0.0675	0.0669

Table 6: Mean absolute changes and standard deviation of changes in exposures to equities (w_{SP}), bonds (w_{TY}) and gold (w_{GC}).

6 Conclusion

Forecasts of volatility are important in many aspects of finance, and as such this literature has grown substantially in recent years. In recent years it has been shown that by harnessing high frequency returns, superior estimates of volatility can be produced. Such estimates, commonly known as realized volatility, are often used for forecasting volatility and lead to superior portfolio allocation outcomes. Traditionally, the parameters of volatility models are estimated within a maximum likelihood framework, and the forecasts they generate are often used, or evaluated in economic applications such as portfolio allocation.

While it is well known that realized volatility is the preferred approach for estimating volatility itself, less is understood about whether the loss function under which model parameters are estimated is important. In this paper we address this issue, and in doing so examine whether realized volatility is still preferred under alternative estimation schemes. We find that while realized volatility based on high frequency data is preferred over simply using daily returns, the loss function under which model parameters are estimated is also important. Using realized volatility, there is little difference in the economic benefit produced under either maximum likelihood or simple mean squared error loss functions. Portfolios generated from forecasts produced under mean squared error are also found to be more stable, an important finding given the impact of transactions costs. From a practical point of view, a means squared error approach also avoids the computation issues of evaluating a multivariate likelihood function when the dimension of the portfolio becomes large.

References

T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, 96(453):42–55, 2001.

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

T.G. Andersen, T. Bollerslev, P.F. Christoffersen, and F.X. Diebold. Volatility and correlation forecasting. In G. Elliot, C.W.J. Granger, and A. Timmerman, editors, *Handbook of Economic Forecasting*. Elsevier, Burlington, 2006.

O.E. Barndorff-Nielsen and N. Shephard. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of Royal Statisitcal Society, Series B*, 64: 253–280, 2002.

R. Becker and A. Clements. Are combination forecasts of s&p 500 volatility statistically superior. *The International Journal of Forecasting*, 24:122–133, 2008.

T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.

T. Bollerslev. Modelling the coherence in short-run nominal exchange rates: A multivariate generalized arch model. *Review of Economics and Statistics*, 72:498–505, 1990.

J. Y. Campbell, A. W. Lo, and A. G. Mackinlay. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, 1997.

M.M. Copeland and T.E. Copeland. Market timing: Style and size rotation using vix. *Financial Analysts Journal*, 55:73–81, 1999.

R.F. Engle. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50:987–1008, 1982.

R.F. Engle. Dynamic conditional correlation. a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics*, 20:339–350, 2002.

J. Fleming, C. Kirby, and B. Ostdiek. The economic value of volatility timing. *The Journal of Finance*, 56:329–352, 2001.

J. Fleming, C. Kirby, and B. Ostdiek. The economic value of volatility timing using realized volatility. *Journal of Financial Economics*, 67:473–509, 2003.

E. Ghysels, P. Santa-Clara, and R. Valkanov. There is a risk-return trade-off after all. *Journal of Financial Economics*, 76:509–548, 2005.

C. Gourieroux and J. Jasiak. *Financial Econometrics*. Princeton University Press, Princeton, New Jersey, 2001.

J.R. Graham and C.R. Harvey. Market timing ability and volatility implied in investment newsletters' asset allocation recommendations. *Journal of Financial Economics*, 42(3):397–421, 1996.

P.R. Hansen and A. Lunde. A forecast comparison of volatility models: Does anything beat a garch(1,1)? *Journal of Applied Econometrics*, 20:873–889, 2005.

S-H. Poon and C. W. J. Granger. Forecasting volatility in financial markets: a review. *Journal of Economic Literature*, 41:478–539, 2003.

S-H. Poon and C. W. J. Granger. Practical issues in forecasting volatility. *Financial Analysts Journal*, 61:45–56, 2005.

G.W. Schwert. Does stock market volatility change over time? *Journal of Finance*, 44:1115–1153, 1989.

S. Skouras. Decisionmetrics: A decision-based approach to economic modelling. *Journal of Econometrics*, 137:414–440, 2007.

K.D. West, H.J. Edison, and D. Cho. A utility-based comparison of some models of exchange rate volatility. *Journal of International Economics*, 35:23–45, 1993.