Clements, Adam & Silvennoinen, Annastiina
(2013)
Volatility timing: How best to forecast portfolio exposures.
Journal of Empirical Finance, *24*, pp. 108-115.

https://doi.org/10.1016/j.jempfin.2013.09.004

Volatility timing: How best to forecast portfolio exposures

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

Abstract

This paper investigates how best to forecast optimal portfolio weights in the context of a volatility timing strategy. It measures the economic value of a number of methods for forming optimal portfolios on the basis of realized volatility. These include the traditional econometric approach of forming portfolios from forecasts of the covariance matrix, and a novel method, where a time series of optimal portfolio weights are constructed from observed realized volatility and directly forecast. The approach proposed here of directly forecasting portfolio weights shows a great deal of merit. Resulting portfolios are of equivalent economic benefit to a number of competing approaches and are more stable across time. These findings have obvious implications for the manner in which volatility timing is undertaken in a portfolio allocation context.

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

JEL Classification Numbers
C22, G11, G17

Acknowledgements The authors would like to thank seminar participants at the Monash University, and in particular Farshid Vahid, and anonymous referees, for their useful comments. Remaining errors are the responsibility of the authors.

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

1 Introduction

The strategy of volatility timing as a portfolio selection method, is often based on forecasts of the volatility of, and correlation between a portfolio's constituent assets. The modern volatility forecasting literature stems from the seminal work of Engle (1982) and Bollerslev (1986) in a univariate setting, and from Bollerslev (1990) and Engle (2002) among others in the multivariate setting. For a broad overview of the major developments in this field, see Gourieroux and Jasiak (2001) and Andersen, Bollerslev, Christoffersen, and Diebold (2006).

A voluminous literature exists dealing with modeling 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 (2005) or Becker and Clements (2008). Relatively speaking, there are fewer studies that focus on the economic value of forecasting volatility, or volatility timing. Graham (1996) 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 has been significant development 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, 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 in the context of volatility timing, relative to estimates based on daily returns.

This paper compares the economic benefit of a range of approaches to volatility timing. Portfolios based on forecasts of volatility from traditional econometric models will be compared to a novel approach where observations of realized volatility are used to construct a time series of optimal portfolio weights, from which forecasts of portfolio weights are directly generated. This method is in contrast with the traditional approach where forecasts of volatility are generated, upon which optimal portfolios are formed.

[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.

The empirical analysis is based on a three asset portfolio allocation problem involving equities, bonds, and gold. The results reveal a number of interesting findings. Forecasting methods that give the greatest weight to the most recent observations, and avoid a great deal of smoothing produce the best performing forecasts. While a very naïve forecast is of similar economic benefit to those that do involve a degree of smoothing, it produces very volatile portfolio exposures. Thus, while some degree of smoothing is beneficial, forecasts involving too much smoothing are inferior. Most interestingly, the novel approach of using observations of realized volatility to construct a time series of optimal portfolio weights, from which forecasts of portfolio weights are formed, is found to be effective. In terms of economic benefit, is it comparable to a traditional forecasting appraoch, but leads to more stable portfolio exposures in a number of instances.

The paper proceeds as follows. Section 2 outlines the general portfolio allocation framework along with how performance will be compared. Details of the competing approaches for forecasting portfolio weights are given in Section 3. Section 4 describes the data employed. Sections 5 and 6 provide empirical results and concluding comments.

2 The portfolio allocation problem and forecast evaluation

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

$$\boldsymbol{r}_t \sim \boldsymbol{F}(\boldsymbol{\mu}, \boldsymbol{\Sigma}_t)\,, \tag{1}$$

where $\boldsymbol{F}$ is some multivariate distribution, $\boldsymbol{\mu}$ is fixed vector of expected returns and $\boldsymbol{\Sigma}_t$ is the conditional covariance matrix of returns.

Following Fleming *et al.* (2003) intraday return information is employed to obtain realized covariances simply denoted here as RV. In this case, the estimate of RV is the sum of the outer-product of intraday returns,

$$\widehat{\boldsymbol{\Sigma}}_t^{RV} = \sum_{i=1}^{N} \boldsymbol{r}_t^i \boldsymbol{r}_t^{i'}, \qquad t = 1, \ldots, T \tag{2}$$

where N represents the number of intraday intervals.

The vector of optimal portfolio weights, $\mathbf{w}_t$, assuming a target portfolio return of μ_0, and proxy for the conditional covariance matrix $\widehat{\boldsymbol{\Sigma}}_t$ is given by

$$\mathbf{w}_t = \frac{\widehat{\boldsymbol{\Sigma}}_t^{-1}\boldsymbol{\mu}}{\boldsymbol{\mu}'\widehat{\boldsymbol{\Sigma}}_t^{-1}\boldsymbol{\mu}}\,\mu_0, \qquad t = 1, \ldots, T. \tag{3}$$

The traditional approach to portfolio allocation is to utilize historical data to generate a forecast of the conditional covariance matrix, $\overline{\boldsymbol{\Sigma}}_{t+1}$. This, in turn, is used to generate a forecast of optimal portfolio weights $\mathbf{w}_{t+1}$ using the relation similar to that of equation (3).

$$\mathbf{w}_{t+1} = \frac{\overline{\boldsymbol{\Sigma}}_{t+1}^{-1}\boldsymbol{\mu}}{\boldsymbol{\mu}'\overline{\boldsymbol{\Sigma}}_{t+1}^{-1}\boldsymbol{\mu}}\,\mu_0\,. \tag{4}$$

The alternative proposed here is to generate a time series of optimal portfolio weights, $\{\mathbf{w}_t\}$, $t = 1, \ldots, T$, based on $\widehat{\boldsymbol{\Sigma}}_t^{RV}$ using intraday returns. A forecast of the optimal portfolio weights $\mathbf{w}_{t+1}$ is then generated directly from the time series of constructed weights $\{\mathbf{w}_t\}$. Both methods will be described in the following section. Irrespective of how the forecast $\mathbf{w}_{t+1}$ is obtained, realized portfolio returns are given by $r_{p,t+1} = \mathbf{w}'_{t+1}\boldsymbol{r}_{t+1}$.

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

$$\sum_{t=T_0}^{T_1} U(r_{p,t}^1) = \sum_{t=T_0}^{T_1} U(r_{p,t}^2 - \delta) \tag{5}$$

where $r_{p,t}^1$ and $r_{p,t}^2$ represent portfolio returns based on two competing forecasting methods, and where T_0 and T_1 date the beginning and the end of the forecasting period, respectively. 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. Following Skouras (2007) an investor with negative exponential utility is considered:

$$U(r_{p,t}) = -\exp(-\lambda\, r_{p,t}) \tag{6}$$

where $r_{p,t}$ is the portfolio return realized by the investor during the period to time t and λ is their coefficient of risk aversion.

The portfolio choice implied by equations (3) and (4) 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 assets are computed from the artificial sample and used as an estimate for $\boldsymbol{\mu}$. Given a series of volatility forecasts, $\{\overline{\boldsymbol{\Sigma}}_{t+1}\}$, portfolio weights are computed using equation (4), and the optimal portfolio returns are recorded. The difference in economic value between any two competing forecasting methods is computed as δ in equation (5). 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 optimal portfolio weights

The methods considered for forecasting optimal portfolio weights fall into two categories. One that forecast the covariance matrix and then subsequently compute the portfolio weights, and one that compute optimal portfolio weights and then directly forecast from these weights. Each approach will now be discussed in turn.

3.1 Forecasting the covariance matrix

The simplest forecasts considered are moving averages of past RV over various horizons, which clearly avoid the need for parameter estimation. Four versions of this simple forecast are employed, the current level of RV, one week and one month moving averages, along an average of the three. The moving averages are defined by

$$\overline{\mathbf{\Sigma}}_{t+1} = \frac{h}{1} \sum_{i=1}^{h} \widehat{\mathbf{\Sigma}}_{t-i+1}^{RV} \tag{7}$$

with h taking the values of $h = 1, 5, 22$ for one day, week and month long moving averages. Subsequently these approaches will be denoted as MA^1, MA^5 and MA^{22}. A simple average of MA^1, MA^5 and MA^{22} will also be considered and will be denoted as MA^{μ}.

Following Fleming *et al.* (2003), the exponentially weighted moving average model employs past RV estimates to generate forecasts, $\overline{\mathbf{\Sigma}}_{t+1}$,

$$\overline{\mathbf{\Sigma}}_{t+1} = \exp(-\theta)\overline{\mathbf{\Sigma}}_t + \theta \exp(-\theta)\widehat{\mathbf{\Sigma}}_t^{RV} \tag{8}$$

where θ is a decay parameter to be estimated. This forecast will be denoted as FKO.

The MIDAS methodology produces volatility forecasts directly from a weighted average of past observations of volatility. Following from Ghysels, Santa-Clara and Valkanov (2006) a forecast of the conditional covariance matrix, $\overline{\mathbf{\Sigma}}_t$ is generated by

$$\overline{\mathbf{\Sigma}}_{t+1} = \sum_{k=1}^{k_{\max}} b\,(k, \boldsymbol{\theta})\, \widehat{\mathbf{\Sigma}}_{t-k+1}^{RV} \tag{9}$$

where the parameters k and $\boldsymbol{\theta} = (\theta_1, \theta_2)'$ govern the MIDAS weighting scheme $b\,(k, \boldsymbol{\theta})$, and $\widehat{\mathbf{\Sigma}}_{t-k+1}^{RV}$ are historical estimates of the realized covariance matrix based on intraday returns. In this instance, the same scalar MIDAS weights, $b\,(k, \boldsymbol{\theta})$ are applied to all elements of the matrix for each lag k. The maximum lag length $k_{\max}$ can be chosen rather liberally as the weights are tightly parameterized. All subsequent analysis is based on $k_{\max} = 100$. The weights are determined by means of a beta density function, fully specified by the parameter vector $\boldsymbol{\theta}$ and normalized such that $\sum b\,(k, \boldsymbol{\theta}) = 1$. To reduce computational burden, and without loss of generality, we restrict $\theta_1 = 1$, which leaves only θ_2 to be estimated. The constraints, $\theta_1 = 1$ and $0 < \theta_2 < 1$, ensure that the weighting function is a decreasing function of the lag k. Both the FKO decay parameter, θ, and the elements in the MID weighting scheme, $\boldsymbol{\theta}$, are determined by standard Quasi-Maximum Likelihood (QML) estimation.

3.2 Directly forecasting portfolio weights

In contrast with the approaches above, we propose forecasting the optimal portfolio weights directly, without resorting to forming a forecast of the covariance matrix. To achieve this, a time series of covariance estimates is formed, with equation (3) used to obtain a series of optimal portfolio weights. The estimates of RV from equation (2) are used to obtain optimal weights $\mathbf{w}_t$, an approach denoted as RVP. Apart from $\widehat{\boldsymbol{\Sigma}}_t^{RV}$, two further estimates of the covariance matrix are used to form $\mathbf{w}_t$, in-sample fitted values, $\overline{\boldsymbol{\Sigma}}_t$ from both the FKO and MID models. Forecasts of portfolio weights determined under these schemes are denoted by RVP^F and RVP^M. This approach is used to determine whether smoothing historical RV prior to determining portfolio weights lead to superior portfolio performance.

The MIDAS scheme can now be used to generate forecasts of the optimal weights directly. Using the MIDAS framework, the relation to be estimated is

$$\mathbf{w}_{t+1} = \sum_{k=1}^{k_{\max}} b\left(k, \boldsymbol{\theta}\right) \mathbf{w}_{t-k+1}. \tag{10}$$

where the historical observations $\mathbf{w}_{t-k+1}$ are generated according to equation (3) with $\widehat{\boldsymbol{\Sigma}}_{t-k+1}^{RV}$ as the covariance proxy. Once again, the same MIDAS weight is applied to all elements in $\mathbf{w}_{t-k+1}$.

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. The estimates of the covariance matrix based on intraday returns were constructed by summing the cross products of 15 minute futures contract returns.

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

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

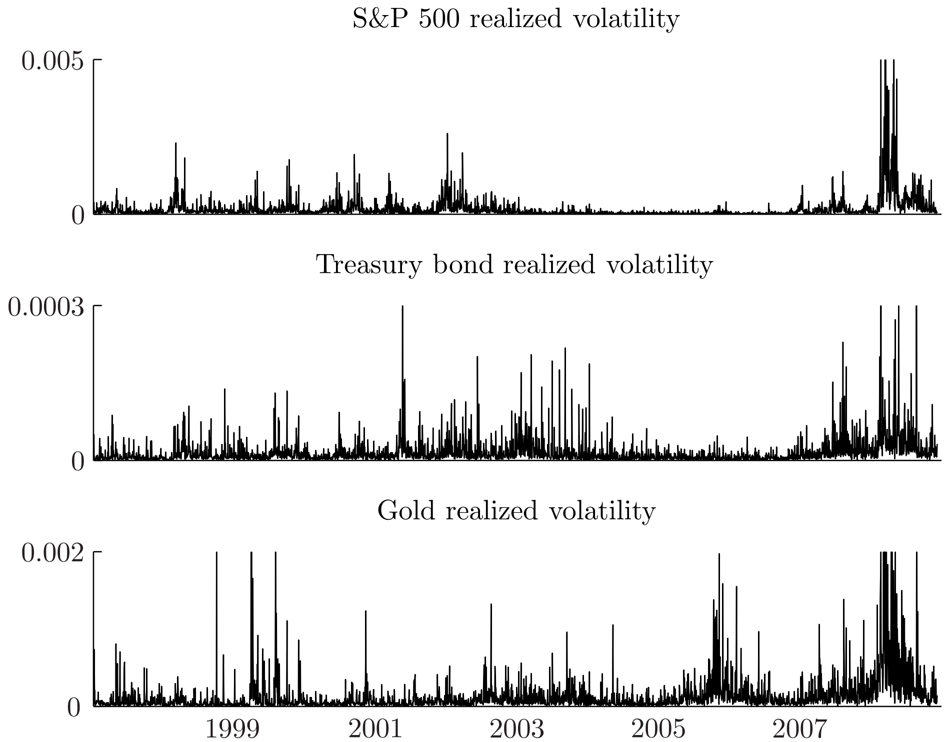

Figure 1: Realized volatility estimates for S&P 500 (top panel), Treasury bond (middle panel) and Gold (bottom panel).

panel) and gold futures returns (bottom 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 increased in late 2005 and early 2006 due to central bank activity, and rose to historically high levels at 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 (bottom panel) show any long-term persistence or structure.

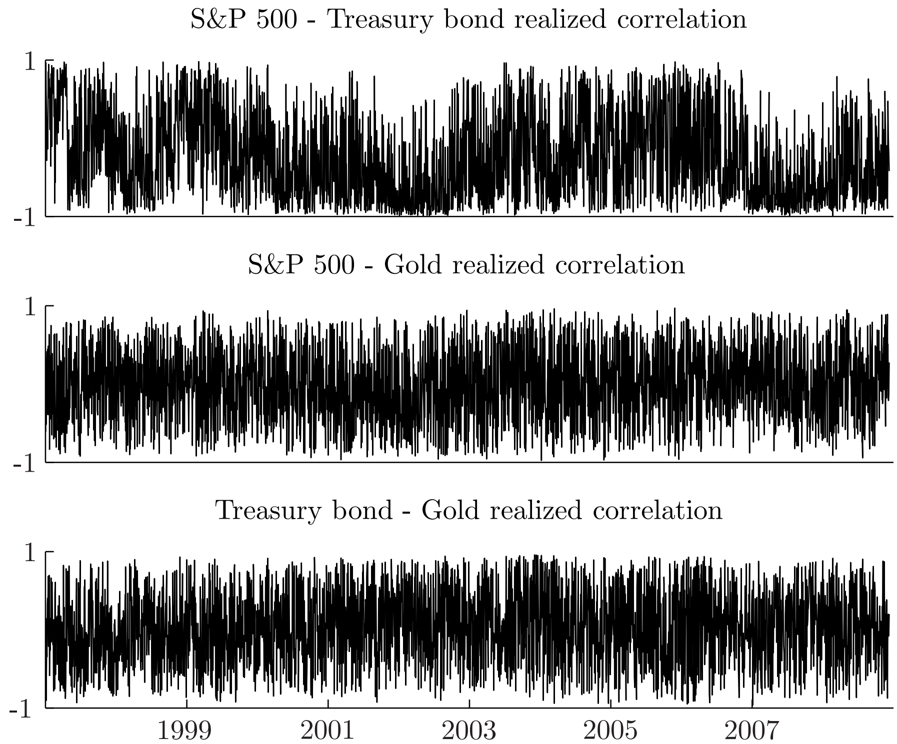

Figure 2: Realized correlation estimates for S&P 500 and Treasury bond (top panel), S&P 500 and Gold (middle panel) and Treasury bond and Gold (bottom panel).

5 Empirical results

Given the 2985 daily observations, the first 1000 observations were used as an initial estimation period. One day ahead forecasts of the optimal weights are obtained from $T = 1000$ onwards and a portfolio formed according to Section 2 which leads to 1985 forecasts of portfolio allocations. All necessary parameters in FKO, MID, RVP, RVP^F and RVP^M are re-estimated every 200 days. As discussed in Section 2, the bootstrap procedure is used to reduce the uncertainty around the expected returns needed in forming the optimal weights. To this end, an artificial sample of returns is generated and the average returns computed to serve as a proxy for the expected returns $\boldsymbol{\mu} = (\mu_{SP}, \mu_{TY}, \mu_{GC})'$. We use the same constraints on the expected returns as Fleming *et al.* (2003). A bootstrap is acceptable if $\mu_{SP} > \mu_{TY} > \mu_{GC}$, $\mu_{SP} > 0$ and $\mu_{TY} > 0$. We also require $\sigma_{SP} > \sigma_{TY}$, where σ denote the sample standard deviation of the artificial returns. Under the FKO and MID approaches, the parameters are not re-estimated for each bootstrap as $\boldsymbol{\mu}$ does not enter the QML objective function during estimation. On the other hand, θ_2 in the MIDAS weighting scheme must be re-estimated under RVP, RVP^F and RVP^M as the optimal weight forecast is a function of past weights which are based on the estimate of $\boldsymbol{\mu}$.

$\mu = 6\%, \gamma = 2$									
	MA^1	MA^5	MA^{22}	MA^{μ}	MID	FKO	RVP	RVP^F	RVP^M
MA^1	-	-87.368	-114.862	-23.079	-1.504	-47.328	-2.302	-50.057	-44.260
		0.084	0.070	0.296	0.426	0.164	0.424	0.188	0.218
MA^5		-	-31.377	39.509	76.072	36.220	71.219	37.555	43.748
			0.082	0.992	1.000	0.988	1.000	0.854	0.892
MA^{22}			-	46.809	97.075	63.926	97.257	72.662	76.903
				0.978	1.000	1.000	0.998	1.000	1.000
MA^{μ}				-	35.046	-5.177	29.440	-4.187	1.256
					1.000	0.118	1.000	0.268	0.442
MID					-	-37.372	-3.009	-29.805	-24.540
						0.000	0.500	0.028	0.032
FKO						-	33.644	9.268	13.804
							0.980	0.830	0.900
RVP							-	-27.777	-22.958
								0.092	0.106
RVP^F								-	4.212
									0.972
RVP^M									-

Table 1: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 6\%$, $\lambda = 2$. Each entry reports the average value for δ across the 500 bootstraps and the proportion of bootstraps where δ was found to be positive.

Tables 1 through 4 report estimates of average δ (averaged across 500 bootstraps) computed from equation (5) for a selection of combinations of target returns μ_0 and degree of risk aversion λ. These represent the incremental economic value of using the forecast from the model in the column heading over that in row heading. These are expressed in annual basis point terms. The number beneath each entry corresponds to the proportion of bootstraps where δ was found to be positive.

To begin, we will focus on Table 1 which reports results for the case of $\mu_0 = 6\%$ and $\lambda = 2$. Of the moving average approaches, MA^1 is the preferred approach as it dominates MA^5, MA^{22} and MA^{μ}. MA^{μ} in fact is preferred over MA^5 and MA^{22} indicating that the information contained in MA^1 is vitally important for forecasting portfolio weights as MA^{μ} is an average of the three and the most recent data dominates. It is seen that MID is of equivalent value to MA^1 and outperforms the other moving average combinations. This result is consistent with the MID approach placing by far the greatest weight on the most recent data as reflected in MA^1. While FKO outperforms MA^5 and MA^{22}, it underperforms both MA^1 and MID indicating that it smooths out too much of the information contained in recent data.

We now consider the performance of the methods that directly forecast portfolio weights. RVP, using only $\widehat{\boldsymbol{\Sigma}}_t$ to form portfolio weights is of equivalent value to MA^1 and also dominates the longer term moving averages. It is also equivalent to MID but dominates FKO. These results are driven by the fact that RVP is directly based on $\widehat{\boldsymbol{\Sigma}}_t$ but the forecast of portfolio weights are smoothed somewhat with the greatest weight placed on the most recent observations, and hence is similar in performance to MA^1 and MID. Finally, the performance of the two methods that utilize a smoothed version of the covariance matrix, $\overline{\boldsymbol{\Sigma}}_t$, are examined. Both RVP^{FKO}

$\mu = 6\%, \gamma = 5$									
	MA^1	MA^5	MA^{22}	MA^{μ}	MID	FKO	RVP	RVP^F	RVP^M
MA^1	-	-56.656 0.158	-88.610 0.108	47.655 0.646	51.726 0.684	-13.574 0.294	42.216 0.674	-64.175 0.156	-52.045 0.186
MA^5		-	-40.248 0.064	41.964 0.976	84.638 1.000	34.889 0.976	65.585 1.000	-4.088 0.436	8.859 0.554
MA^{22}			-	20.632 0.908	98.082 1.000	65.684 1.000	91.660 0.998	45.757 1.000	53.671 1.000
MA^{μ}				-	39.045 1.000	-11.653 0.066	17.961 0.870	-51.663 0.030	-40.577 0.034
MID					-	-43.285 0.000	-14.545 0.018	-66.447 0.000	-55.923 0.000
FKO						-	26.833 0.940	-18.638 0.080	-9.984 0.256
RVP							-	-54.235 0.030	-44.831 0.038
RVP^F								-	7.772 0.996
RVP^M									-

Table 2: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 6\%$, $\lambda = 5$. Each entry reports the average value for δ across the 500 bootstraps and the proportion of bootstraps where δ was found to be positive.

and RVP^{MID} are inferior MA^1, superior to MA^5 and MA^{22} and similar in performance to MA^{μ}. They are clearly inferior to both MID and RVP, though marginally dominate FKO. Overall RVP^{FKO} and RVP^{MID} are very similar in nature. As both approaches are inferior to the dominant methods, MA^1, MID and RVP, smoothing the covariance estimates prior to forming portfolios reduces the amount of information reflected in recent covariances, and or portfolio weights relevant for forecasting purposes.

The results in Table 3 show that increasing the target return from $\mu_0 = 6\%$ to $\mu_0 = 8\%$ has no impact on the order of preference between the competing forecasting approaches. Overall, the incremental economic benefit, δ does seem to strengthen in magnitude making the differences between model performance even clearer. MA^1, MID and RVP continue to perform in a similar manner, with the proportion of times that MID or RVP dominate MA^1 range from 40% to 70% of the bootstraps highlighting no significant differences. These approaches continue to dominate the other methods. A similar conclusion can be drawn from results in Tables 2 and 4 when increasing the coefficient of risk aversion. Only in the case of $\mu_0 = 8\%$ and $\gamma = 5$ is there is a significant difference in the performance where RVP slightly dominates MID. Overall, the result that MA^1, MID and RVP are dominant methods indicate that methods that place a great deal of weight on the most recent observations (the most recent observation in the case of MA^1) lead to superior performance. By smoothing the covariance estimates first, much of the information pertaining to the forecasts of optimal weights is lost.

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

$\mu = 8\%$, $\gamma = 2$									
	MA^1	MA^5	MA^{22}	MA^{μ}	MID	FKO	RVP	RVP^F	RVP^M
MA^1	-	-107.231 0.108	-144.587 0.074	-9.858 0.380	14.145 0.472	-52.432 0.194	10.710 0.474	-69.552 0.182	-60.075 0.210
MA^5		-	-44.052 0.078	53.268 0.986	104.078 1.000	48.206 0.990	93.525 1.000	38.531 0.796	48.687 0.832
MA^{22}			-	54.342 0.976	129.567 1.000	85.693 1.000	127.963 0.998	89.241 1.000	95.906 1.000
MA^{μ}				-	48.145 1.000	-8.385 0.112	36.234 0.994	-18.697 0.118	-9.875 0.240
MID					-	-51.427 0.000	-7.295 0.388	-49.994 0.012	-41.502 0.022
FKO						-	42.838 0.970	4.492 0.654	11.680 0.786
RVP							-	-44.432 0.064	-36.732 0.082
RVP^F								-	6.606 0.990
RVP^M									-

Table 3: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 8\%$, $\lambda = 2$. Each entry reports the average value for δ across the 500 bootstraps and the proportion of bootstraps where δ was found to be positive.

$\mu = 8\%$, $\gamma = 5$									
	MA^1	MA^5	MA^{22}	MA^{μ}	MID	FKO	RVP	RVP^F	RVP^M
MA^1	-	-53.396 0.202	-102.906 0.116	116.323 0.852	106.546 0.882	3.362 0.370	86.235 0.842	-103.793 0.136	-82.116 0.166
MA^5		-	-62.643 0.042	58.335 0.972	118.541 1.000	43.732 0.948	81.884 1.000	-41.129 0.230	-18.209 0.356
MA^{22}			-	9.538 0.710	132.441 1.000	89.241 1.000	118.414 0.996	39.072 0.964	52.835 0.990
MA^{μ}				-	53.829 1.000	-22.720 0.050	13.378 0.786	-109.638 0.014	-90.002 0.018
MID					-	-62.975 0.000	-28.682 0.000	-119.375 0.000	-100.882 0.000
FKO						-	30.775 0.918	-47.971 0.006	-32.873 0.038
RVP							-	-94.617 0.020	-78.149 0.026
RVP^F								-	13.406 1.000
RVP^M									-

Table 4: Estimates of relative economic value of moving from the forecasts in the row heading to that in the column heading, for $\mu_0 = 8\%$, $\lambda = 5$. Each entry reports the average value for δ across the 500 bootstraps and the proportion of bootstraps where δ was found to be positive.

	MA^1	MA^5	MA^{22}	MA^{μ}	MID	FKO	RVP	RVP^F	RVP^M
$\|\Delta_{w_E}\|$	0.2886	0.0714	0.0153	0.0937	0.0271	0.0106	0.0287	0.0031	0.0049
$\sigma_{\Delta_{w_E}}$	0.2983	0.0900	0.0197	0.1015	0.0332	0.0151	0.0303	0.0033	0.0061
$\|\Delta_{w_B}\|$	0.4546	0.1149	0.0243	0.1507	0.0440	0.0171	0.0450	0.0049	0.0080
$\sigma_{\Delta_{w_B}}$	0.4566	0.1392	0.0300	0.1580	0.0515	0.0231	0.0468	0.0052	0.0095
$\|\Delta_{w_G}\|$	0.2989	0.0710	0.0152	0.0932	0.0289	0.0109	0.0292	0.0029	0.0051
$\sigma_{\Delta_{w_G}}$	0.3404	0.0831	0.0183	0.0997	0.0345	0.0154	0.0350	0.0034	0.0064

Table 5: Mean absolute changes and standard deviation of changes in exposures to equities (w_E), bonds (w_B) and gold (w_G).

absolute changes in portfolio weights and their standard deviations, to reveal whether a link exists between the competing forecasts and portfolio stability. The mean absolute changes in exposures to a single asset class, equities bonds or gold, are defined as the mean absolute change in exposure averaged across the N number of bootstraps,

$$|\Delta_w| = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{T_1 - T_0 - 1}\sum_{t=T_0+1}^{T_1}|w_t - w_{t-1}| \tag{11}$$

where T_0 and T_1 are the first and final forecast periods. The standard deviation of exposure changes is determined in a similar manner,

$$\sigma_{\Delta_w} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{T_1 - T_0 - 1}\sigma_{w_t - w_{t-1},i} \tag{12}$$

where $\sigma_{w_t - w_{t-1},i}$ is the standard deviation of portfolio changes for the ith bootstrap. These statistics are reported in Table 5 for each of the three asset classes. It is clear that using the MA^1 covariance forecast leads to the most variable portfolio exposures, with MA^5 and MA^{22} being noticeably more stable, an expected result as MA^1 involves no smoothing or averaging. Exposures resulting from MA^{22} are more volatile again as more weight is placed on the more variable MA^1 forecast. Variability in exposures from MID and RVP are broadly similar with FKO producing slightly less variability. While the mean absolute changes for RVP are slightly larger than MID, the volatility of the exposure changes from RVP are nearly 10% lower in the equity and bond cases. Finally, by first smoothing the covariance estimates, RVP^{FKO} and RVP^{MID} produce the most stable portfolio exposures.

Overall these results indicate that forecasting methods that give the greatest weight to the most recent observations, and avoid a great deal of smoothing produce the best performing forecasts. While the naïve forecast MA^1 is of similar economic benefit to those that do involve a degree of smoothing, MID and RVP, it leads to portfolio exposures that are much too volatile. Thus, while some degree of smoothing is beneficial, forecasts involving too much smoothing are inferior in terms of economic benefit. The novel approach of computing optimal portfolio weights and directly forecasting the exposures has some merit, it is of equal economic benefit to the traditional forecasting approach of MID but leads to a reduction in the volatility of exposures. Potentially, the transformation to optimal portfolio weights may reduce estimation error as fewer elements must be forecast moving from the full covariance matrix to the portfolio weights.

6 Conclusion

The issue of volatility timing has attracted a great deal of research attention. Traditionally, an econometric model of volatility is used to generate covariance forecasts upon which optimal portfolios are based. There is little understanding of how best to estimate such models for the purposes of volatility timing. Therefore, this paper undertakes a comparison of methods for forecasting optimal portfolio weights in the context of volatility timing. A novel approach to volatility timing is also proposed. It involves constructing a time series of observed optimal portfolio weights, upon which forecasts are based. Results indicate that while a naïve forecast produces forecast of similar economic benefit to competing time-series forecasts, it leads to very volatile portfolio exposures. The proposed approach of forecasting portfolio weights directly appears to have merit in that it leads to portfolios of equivalent economic benefit to the traditional approach of forecasting the covariance matrix and leads to more stable portfolio exposures. This result opens a potentially new avenue for research that offers an alternative to the use of traditional economic models of volatility.

References

Andersen, T.G., Bollerslev, T., Christoffersen, P.F. and Diebold, F.X. (2006), Volatility and Correlation Forecasting, in *Handbook of Economic Forecasting*, eds. Elliot, G., Granger, C.W.J. and Timmerman, A., Elsevier, Burlington.

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

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

Becker, R., and Clements, A. (2008), "Are combination forecasts of S&P 500 volatility statistically superior?", *The International Journal of Forecasting*, **24**, 122-133.

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

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

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

Copeland, M.M. and Copeland, T.E. (1999), "Market timing: Style and size rotation using VIX ", *Financial Analysts Journal*, **55**, 73-81.

Engle, R.F. (1982), "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation", *Econometrica,* **50**, 987-1007.

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

Fleming, J., C. Kirby and B. Ostdiek. (2001), "The economic value of volatility timing", *The Journal of Finance*, **56**, 329-352.

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

Ghysels, E., Santa-Clara, P., Valkanov, R., 2006. Predicting volatility: How to get most out of returns data sampled at different frequencies. Journal of Econometrics, 131, 59-95.

Gourieroux, C. and Jasiak, J. (2001), *Financial Econometrics*. Princeton University Press, Princeton NJ.

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

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

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

Poon, S-H. and Granger, C.W.J. (2005), ”Practical Issues in forecasting volatility”, *Financial Analysts Journal*, **61**, 45-56.

Schwert, G.W. (1989), “Does stock market volatility change over time?”, *Journal of Finance*, **44**, 1115-1153.

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

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