

On the economic benefit of utility based estimation of a volatility model

A Clements and A Silvennoinen

School of Economics and Finance, Queensland University of Technology.

Abstract

Forecasts of asset return volatility are necessary for many financial applications, including portfolio allocation. Traditionally, the parameters of econometric models used to generate volatility forecasts are estimated in a statistical setting and subsequently used in an economic setting such as portfolio allocation. Differences in the criteria under which the model is estimated and applied may inhibit reduce the overall economic benefit of a model in the context of portfolio allocation. This paper investigates the economic benefit of direct utility based estimation of the parameters of a volatility model and allows for practical issues such as transactions costs to be incorporated within the estimation scheme. In doing so, we compare the benefits stemming from various estimators of historical volatility in the context of portfolio allocation. It is found that maximal utility based estimation, taking into account transactions costs, of a simple volatility model is preferred on the basis of greater realized utility. Estimation of models using historical daily returns is preferred over historical realized volatility.

Keywords

Volatility, utility, portfolio allocation, realized volatility, MIDAS

JEL Classification Numbers C10, C22, G11, G17

Pick some codes

Corresponding author

Adam Clements

School of Economics and Finance

Queensland University of Technology

Brisbane, 4001

Qld, Australia

email a.clements@qut.edu.au

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 (1986). For broad overview of the major developments in this field, see for instance Campbell, Lo and MacKinlay (1997), Gouriéroux and Jasiak (2001) and Andersen, Bollerslev, Christoffersen and Diebold (2006).

Traditionally, volatility models are estimated in a statistical setting, most commonly embedding a parametric volatility model into a conditional probability density function. This approach has been used in the context of the family of GARCH models, and it has been refined by replacing the Gaussian density function with one that better suits the excess kurtosis and, or skewness (constant or time-varying) left in the standardized residuals. The maximum (quasi) likelihood estimators then provide parameter estimates for the model and possibly for the likelihood function as well that best link the observed data to the assumed data generating process. More recently, another widely used family of models are the MIDAS regressions (Ghysels, Santa-Clara and Valkanov, 2006), in which the structure of the weighting scheme is linked to a density function.

The forecasts generated from estimated volatility models are often harnessed within an economic application such as portfolio allocation. The forecast performance of volatility models are usually evaluated by the use of measures of statistical accuracy such as mean squared error, or economic benefit given a particular economic application such as value-at-risk or portfolio allocation. Poon and Granger (2003, 2005) and provide a wide ranging review of forecast performance from a statistical viewpoint, whereas Andersen, Bollerslev, Christoffersen and Diebold (2006) also discuss a number of economic settings.

Given that likelihood and economic benefit based loss functions will differ, it is not immediately obvious that a model selected under the former will be optimal under the latter. However, the ultimate test for any volatility model should be its performance within the context of the application at hand, and hence the model estimation and selection should be considered together with the end goal. Therefore, it is of interest to compare different estimation approaches in an economic application, such as portfolio allocation. Skouras (2007) proposes a utility based approach for estimating the parameters of a volatility model. Thus, the criteria for estimating the model parameters and the final application of the model are consistent.

Within a portfolio allocation framework, this paper compares the performance of a simple volatility model estimated by traditional statistical method and utility based criterion. The

model chosen is the MIDAS approach of Ghysels *et al.* (2006) with forecasts being based on past volatility measured by daily squared returns, and realized volatility developed by Andersen, Bollerslev, Diebold and Labys (2001, 2003). The results from this comparison will lead to the following contributions. First, the paper will examine the economic benefit of selecting a volatility model through a utility based criterion as opposed to a statistical one. Second, in doing so we will examine the benefits of using intraday data instead of daily data, both when using statistical and utility based estimation criteria. This will reveal interesting interactions between the criteria used to estimate the model and the benefit of using intraday data for measuring volatility. Third, we will also investigate the gains of directly incorporating transaction costs into the estimation procedure. These comparisons give practitioners an idea of how their choice of model selection and its potential link to the economic application can affect the end result. The paper proceeds as follows. Section 2 outlines the data utilized here. Section 3 outlines the volatility model considered, a utility based estimation approach, how transactions costs may be incorporated along with how model performance will be compared. Sections 4 and 5 provide empirical results and concluding comments respectively.

2 Data

The portfolio allocation problem considered here relates to a mix of a single risky (equities) asset and riskless asset. The study treats the S&P 500 Composite Index as the risky asset, with data gathered over the period 2 January 1990 to 31 December 2008 (4791 daily observations).

Daily returns on the physical index itself have been gathered along with daily returns on the futures contracts written on the index. The latter represent excess returns over the riskless rate of return as futures contracts are not cash collateralized. Futures returns are based on a continuous series of futures prices where contracts are rolled to the next nearest maturity contract a week prior to the maturity of the shorter dated contract.

To examine the benefit of using intraday data to form estimates of volatility, the realized volatility (RV) methodology outlined in ABDL (2001, 2003) is used. RV estimates volatility by means of aggregating intra-day squared returns. It should be noted that the daily trading period of the S&P500 is 6.5 hours and that overnight returns were used as the first intra-day return in order to capture the variation over the full calendar day. ABDL (1999) suggest how to deal with practical issues relating to intra-day seasonality and sampling frequency when dealing with intra-day data. Based on the volatility signature plot methodology, daily RV estimates were constructed using 30 minute S&P500 index returns¹. It is widely acknowledged that RV is a

¹Intraday S&P 500 index data were purchased from Tick Data, Inc.

more accurate and less noisy estimate of the unobservable volatility process than squared daily returns, Poon and Granger (2003). Patton 2006 suggests that this property of RV is beneficial when RV is used a proxy for observed volatility when evaluating forecasts.




Figure 1: Daily S&P 500 index returns (top panel) and daily RV estimates (bottom panel).

Figure 1 plots both the daily returns on the S&P index and associated RV estimates (shown in form of daily variances). Both series show a familiar pattern, low volatility during much of the 1990's with periods of higher volatility due to events such as the Asian crisis and collapse of technology stocks interspersed. 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.

3 Methodology

This section will outline the methodology employed in this article. It will begin with a brief description of the simple volatility model based on the principle of MIDAS regressions. The portfolio allocation based economic criteria used for estimating the model parameters is then described, followed by how the economic benefit of the competing approaches will be compared.

3.1 A simple volatility model

The volatility model utilized here is from the family of MIDAS regressions. This methodology produces volatility forecasts directly from a weighted average of past observations of volatility.

This approach is chosen as it is simple to estimate and there is a clear link between the estimated parameters and manner in which historical data is weighted.

Following the notation introduced in Ghysels *et al.* (2006) a forecast is generated according to

$$Q_{t+1} = \sum_{k=0}^{k_{\max}} b(k, \boldsymbol{\theta}) x_{t-k} + \varepsilon_t \quad (1)$$

where Q_{t+1} represents the volatility forecast at time t for periods $t + 1$ conditional upon information available at time t .² Any variable containing information regarding the volatility process can be substituted for x_t . Examples used in Ghysels *et al.* (2006) were RV, squared daily return, absolute daily return, daily range and realized power variation. In this paper we will utilize historical observations of both RV and squared daily returns. Under statistical estimation, Q_{t+1} will be defined as a forecast of variance, $\hat{\sigma}_{t+1}^2$. Using utility based estimation, Q_{t+1} represents the optimal allocation to equity futures. The maximum lag length k_{\max} can be chosen rather liberally as the weight parameters $b(k, \boldsymbol{\theta})$ 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 $\sum b(k, \boldsymbol{\theta}) = 1$. 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 .

Within the statistical context, the parameter θ_2 will be estimated via two methods. When the model is estimated on the basis of RV, values for both x_t and Q_{t+1} are simply RV_t and RV_{t+1} . Estimation of θ_2 is achieved by nonlinear least squares by minimizing the sum of squared residuals in equation (1). Forecasts from this approach will be denoted by RV^s below. When values for x_t in equation (1) are taken to be daily squared returns, θ_2 is estimated using maximum likelihood in the same manner as a GARCH style model. Q_{t+1} is a forecast of variance, $\hat{\sigma}_{t+1}^2$ and ε_t in equation (1) is irrelevant. This model will be denoted as Day^s .

Figure 2 shows the weighting function applied to historical data for a range of parameters, $\theta_1 = 1$ and $\theta_2 = 0.9, 0.5, 0.1$. When $\theta_2 = 0.9$ there is little decay in the weights and hence this weighting function is similar in nature to a simple moving average. However, as the value for θ_2 becomes smaller, more weight is placed on the most recent observations and less on the distant observations.

3.2 The portfolio allocation problem

Following Skouras (2007), we consider an investor with negative exponential utility,

$$u(w) = -\exp(-\lambda w) \quad (2)$$

²MIDAS models are more general as indicated by the notation here. They can deal with data being sampled at different frequencies and can also directly utilize intra-day data. These generalizations are not required here.




Figure 2: MIDAS weighting functions for $\theta_1 = 1$ and $\theta_2 = 0.9, 0.5, 0.01$.

where w is the investor's level of wealth and λ is their coefficient of risk aversion. Denoting μ as expected excess returns, and assuming that returns are distributed $N(\mu, \sigma_{t+1}^2)$ means the optimal weight to hold in risky equities is given by

$$\alpha_{t+1} = \frac{1}{\lambda} \frac{\mu}{\hat{\sigma}_{t+1}^2}. \quad (3)$$

Given this optimal portfolio choice, utility realized by the investor is

$$u(r_{t+1}) = -\exp(-\alpha_{t+1} r_{t+1}) \quad (4)$$

where r_{t+1} is a realization of the excess returns on equities. Given the forecasts, $\hat{\sigma}_{t+1}^2$ obtained from the models estimated using statistical criteria, RV^s and Day^s , the allocation to equities can be obtained from equation (3). Once the portfolio allocation has been established the economic value of the statistical based forecasts can be derived using equation (4). This economic criteria of course differs from the likelihood loss function used to estimate the model parameters using the statistical criteria. The following section will describe how we estimate the volatility model, θ_2 specifically, directly employing the economic framework described here in this section.

3.3 Utility based estimation of the volatility model

Skouras (2007) proposes a method by which the parameters of a 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 (4), the objective function for a maximum utility estimator is

$$\arg \max_{\theta_2} \frac{1}{T} \sum_{t=1}^T -\exp(-\alpha_{t+1} r_{t+1}) = \arg \max_{\theta_2} \frac{1}{T} \sum_{t=1}^T -\exp\left(-\frac{\mu}{\widehat{\sigma}_{t+1}^2} r_{t+1}\right). \quad (5)$$

Once again the MIDAS weighting described in section 3.1 will be used to generate the forecasts, $\widehat{\sigma}_{t+1}^2$ required in the objective function of equation (5). Models will be once again estimated using both historical RV and daily squared returns and will be denoted below as RV^u and Day^u respectively.

3.4 Transaction costs

In practice, an investor will incur transaction costs as they alter their portfolio as a result of changes in the optimal portfolio allocation to equity futures stemming from equation (3). All subsequent empirical analysis is based on the S&P 500 futures contract, highly liquid contract. While this may be the case, an investor will still experience costs from trading, which we assume here to reflect the bid-ask spread. This cost is approximated by $\frac{\text{Bid-ask spread}}{\text{Futures Price}}$ where the bid-ask spread quoted in index points. This is the case as both the bid-ask spread and the value of one futures contract quoted in dollar terms would both be multiplied by the same multiplier, currently \$250 per index point.

We assume this cost is paid by the investor as the optimal allocation for equation (3) changes through time. Assuming an arbitrarily large investment portfolio (or infinitely divisible contracts) the transaction costs are given by

$$tc_{t+1} = |\alpha_{t+1} - \alpha_t| \frac{\text{Bid-ask spread}}{\text{Futures Price}}. \quad (6)$$

This scheme is applied across all models, RV^s , Day^s , RV^u and Day^u by augmenting the expression for realized utility in equation (4),

$$u(r_{t+1}) = -\exp(-(\alpha_{t+1} r_{t+1} - tc_{t+1})). \quad (7)$$

This provides a post-transactions cost measure of the economic value of each of the forecasting models.

An alternative approach for incorporating transactions costs into the utility based models, RV^u and Day^u is to include a penalty term directly into the objective function in equation (5). This leads to an objective function of the following form,

$$\arg \max_{\theta_2} \frac{1}{T} \sum_{t=1}^T -\exp\left(-\left(\frac{\mu}{\widehat{\sigma}_{t+1}^2} r_{t+1} - tc_{t+1}\right)\right) \quad (8)$$

which allows for the impact of transactions costs to influence the estimation of the volatility model. This approach will be used with both historical RV and daily squared returns, with these models denoted below as RV^{utc} and Day^{utc} .

3.5 Comparing economic value across forecasting models

We follow Fleming, Kirby and Ostdiek (2001, 2003) to compare the performance of the various methods for portfolio formation, RV^s , Day^s , RV^u , Day^u , RV^{utc} and Day^{utc} . We find a constant, Δ solving

$$\sum_{t=1}^T u(r_{t,1}) = \sum_{t=1}^T u(r_{t,2} - \Delta) \quad (9)$$

where $r_{t,1}$ and $r_{t,2}$ represent portfolio returns based on two competing forecasting methods. Here Δ reflects the incremental value of using the second model 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 model. Δ will be reported in annualized basis points below.

4 Empirical results

Figure 3 plots the portfolio allocation to equities implied by the RV^s and Day^s (top panel) and RV^u and Day^u (bottom panel) for each of the 3791 one-day ahead forecasts. Allocations shown here relate to expected returns of 4% p.a, $\gamma = 5$ and transactions cost of a 1 index point bid-ask spread³. When comparing the behavior of the statistical based allocation in the top panel an interesting pattern emerges. During periods when allocations to equities are rising (volatility forecasts falling) the allocations given Day^s rise slower than those due to RV^s . This may be due to the fact that daily squared returns are a more noisy proxy for volatility than RV and hence it is harder to distinguish when volatility is falling. On the other hand when volatility is rising, and hence allocations falling, there is little to distinguish between the RV^s and Day^s . The pattern to emerge when comparing the allocations of the utility based models, RV^u and Day^u in the lower panel of Figure 3 is quite different. In this case, the allocations implied by RV^u rise more slowly in a similar manner to Day^u (in turn similar to Day^s). As these forecasts are based on maximizing realized utility, they will reflect the higher penalty associated with under-prediction of volatility. This is consistent with the differences observed between the allocations from RV^s and RV^u . Allocations generated by RV^{utc} and Day^{utc} are not shown here as they are qualitatively similar to those of RV^u and Day^u .

³Differences in allocations across the models are qualitatively similar for other combinations of expected return, risk aversion and transaction costs.




Figure 3: Portfolio allocations to equities. Allocations implied by statistical based forecasts, RV^s and Day^s (top panel) and utility based estimates, RV^u and Day^u (bottom panel).

Estimates of θ_2 and hence the MIDAS weighting function also vary across the models. Estimates of θ_2 are smallest for RV^s with it never reaching 0.05 in any of the estimation windows. Given the weighting functions shown in Figure 2, this implies that the greatest weight is given to very recent observations and little to more distant data. On the other hand, the utility based models, RV^u and Day^u lead to estimates of θ_2 rarely below 0.5 and above 0.9 for many of the estimation windows. This results indicates that optimal allocations based on a maximal utility criterion place a more even weight across historical data.

	RV^s	Day^s	RV^u	Day^u	RV^{utc}	Day^{utc}
RV^s	0.000	29.851	20.731	32.041	31.552	46.029
Day^s		0.000	-9.120	2.190	1.701	16.177
RV^u			0.000	11.310	10.821	25.298
Day^u				0.000	-0.489	13.988
RV^{utc}					0.000	14.477
Day^{utc}						0.000

Table 1: The value of using models in column headings instead of models in row headings, expressed in annualized basis points. The level of expected returns is 4% p.a., the coefficient of risk aversion is $\gamma=5$, and transactions costs are a bid-ask spread of 1 index point.

Tables 1 through 4 reports differences in the economic value of the models as captured in annualized basis point fees for switching between the models. These values reflect the maximum return paid by an investor for switching from models in the row headings of each table to those

	RV ^s	Day ^s	RV ^u	Day ^u	RV ^{utc}	Day ^{utc}
RV ^s	0.000	43.577	33.814	40.196	49.644	62.319
Day ^s		0.000	-9.763	-3.381	6.067	18.742
RV ^u			0.000	6.382	15.831	28.506
Day ^u				0.000	9.449	22.123
RV ^{utc}					0.000	12.675
Day ^{utc}						0.000

Table 2: The value of using models in column headings instead of models in row headings, expressed in annualized basis points. The level of expected returns is 4% p.a., the coefficient of risk aversion is $\gamma=5$, and transactions costs are a bid-ask spread of 2 index points.

	RV ^s	Day ^s	RV ^u	Day ^u	RV ^{utc}	Day ^{utc}
RV ^s	0.000	52.099	35.151	63.967	45.088	75.121
Day ^s		0.000	-16.948	11.869	-7.011	23.022
RV ^u			0.000	28.816	9.937	39.970
Day ^u				0.000	-18.879	11.154
RV ^{utc}					0.000	30.033
Day ^{utc}						0.000

Table 3: The value of using models in column headings instead of models in row headings, expressed in annualized basis points. The level of expected returns is 6% p.a., the coefficient of risk aversion is $\gamma=5$, and transactions costs are a bid-ask spread of 1 index point.

	RV ^s	Day ^s	RV ^u	Day ^u	RV ^{utc}	Day ^{utc}
RV ^s	0.000	72.695	55.360	81.984	72.080	99.450
Day ^s		0.000	-17.335	9.290	-0.615	26.755
RV ^u			0.000	26.624	16.720	44.090
Day ^u				0.000	-9.904	17.466
RV ^{utc}					0.000	27.370
Day ^{utc}						0.000

Table 4: The value of using models in column headings instead of models in row headings, expressed in annualized basis points. The level of expected returns is 6% p.a., the coefficient of risk aversion is $\gamma=5$, and transactions costs are a bid-ask spread of 2 index point2.

in the column headings and are obtained from equation (9). Positive (negative) values indicate that an investor would prefer the model in the column (row) heading. Results are reported for expected returns of 4% p.a. and 6% p.a., $\gamma = 5$ and bid-ask spread transactions costs of 1 and 2 index points⁴.

From the results in Tables 1 through 4 a number of interesting patterns emerge. Of the statistical based models, Day^s is preferred to RV^s . The differences in performance are larger for expected returns of 6% p.a. and transactions costs of 2 index points, relative to expected returns of 4% p.a. and transactions costs of 1 index point. Overall, allocations from utility based estimation are preferred to those based on statistical estimation. Both RV^u and Day^u are preferred to RV^s and Day^s respectively. Once again, the differences in performance are greater for larger transactions costs and expected returns. Of the utility based models, Day^u is preferred to RV^u . Finally, incorporating transaction costs directly into the estimation of model parameters is beneficial, as Day^{uttc} is consistently preferred to all competing models.

In summary, these results indicate that utility based estimation is preferred to the statistical estimation of models for the purposes of portfolio allocation. From the practical point of view this reflects the fact that the criteria under which the model is estimated is consistent with that under which it is applied. Therefore, an inconsistency between the criteria used for estimation and evaluation seems to be costly. The use of daily returns and the incorporation of transactions costs directly into the estimation of models leads to the most beneficial outcomes. Historical data is weighted relatively evenly in the context of utility based estimation, implying that a weighting similar to a simple moving average is acceptable.

5 Conclusion

Forecasts of volatility are important in many aspects of finance, and as such this literature has grown substantially in recent years. While volatility models are traditionally estimated within a statistical framework, the forecasts they generate are often used or evaluated in economic applications such as portfolio allocation. While this is the case, there is little understanding of the differences between models estimated on either a statistical, or directly on economic basis and finally applied to a portfolio allocation problem. This paper seeks to gain a deeper understand how models estimated under both statistical and criteria perform in the portfolio allocation setting.

Within an negative exponential utility framework, it is found that maximal utility based estima-

⁴Results have been generated for values of $\gamma = 2, 10$ but are not reported for the sake of brevity. The results do not differ qualitatively from those reported here and are available upon request

tion of a simple volatility model is preferred on the basis of greater realized utility. As opposed to a statistical forecast of volatility being used to form optimal portfolios, an investor would prefer to estimate a model of optimal allocations as a direct function of historical data. In doing so, the use of daily returns (as opposed to estimates of realized volatility) and the incorporation of transactions costs directly into the estimation of models leads to the best performing portfolios.

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. & Labys, P. (1999). (Understanding, optimizing, using and forecasting) Realized Volatility and Correlation. Working Paper, University of Pennsylvania.
- 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.
- Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, **31**, 307-327.
- Campbell, J.Y., Lo, A.W. and MacKinlay, A.G. (1997), *The Econometrics of Financial Markets*, Princeton University Press, Princeton NJ.
- Engle, R.F. (1982), "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation", *Econometrica*, **50**, 987-1007.
- 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.
- Gourieroux, C. and Jasiak, J. (2001), *Financial Econometrics*. Princeton University Press, Princeton NJ.
- Engle, R. and Colacito, R. (2006), "Testing and valuing dynamic correlations for asset allocation ", *Journal of Business and Economic Statistics*, **24**, 238-253.
- 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.
- González-Rivera, G., Lee, T. and Mishra, S. (2003), "Forecasting volatility: A reality check based on option pricing, utility function, value-at-risk and predictive likelihood ", unpublished manuscript.

Patton, A. (2006), Volatility Forecast Comparison Using Imperfect Volatility Proxies, Quantitative Finance Research Centre, University of Technology, Sydney, research Paper Series, n. 175.

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.

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.