A Comparison between Alternative Volatility Estimations

Application on Blood Oxygen Concentration of Preterm Infants

Zhao, X. [1] , Q.Hou [1] , D.Lee [1], M.Reale [1] , C.Scarrott [1] ,G.Russell [2] , A. MacDonald [1], M.Zahari [1]

[1] Department of Mathematics and Statistics, University of Canterbury, Christchurch

[2] Christchurch Women's Hospital, Christchurch

Email: xzh77@student.canterbury.ac.nz

Keywords: *Stochastic volatility, realized volatility, particle filter*

EXTENDED ABSTRCT

The medical measurements for preterm babies manifest instabilities because of their under-developed biological systems. These physiological instabilities raise the question of whether a preterm baby is simply premature or is likely to develop long-term health complications. Because instabilities are of concern, useful information about the health state of preterm infants can be expected to be contained in the variations of the medical measurements. However, instantaneous variations are not observed directly along with the physiological measurements, and so suitable models must be employed to estimate them.

Volatility models derived from the finance area are used to find the underlying variations for blood oxygen concentration of preterm babies. Although volatility models are not commonly used in medical applications, they have been widely discussed in the finance literature, for example, models such as ARCH, GARCH and SVM. This paper considers two latent volatility models - a Stochastic Volatility Model (SVM) fitted using Bayesian inference and a particle filter, and the EGARCH model. An alternative Realized Volatility estimator, which is recently widely discussed in finance area and regarded as an approximation of realized variance based on high frequency intra-period data, is used as a benchmark to compare the two latent volatility estimators.

The relative performance of the two latent volatility models is first evaluated using the R^2 measure from a linear regression analysis involving realized volatility and each of the latent volatility models separately. Several other numerical measures of closeness between the latent volatility estimates and realized volatility are also computed. They all show the SVM estimates to be closer to realized volatility than the EGARCH estimates. Finally, the bootstrap method is used to obtain point-wise confidence intervals for realized volatility, to further assess the latent volatility estimates.

The empirical results show that, consistent with return volatility in finance, the volatilities of the oxygen level of preterm babies, estimated by stochastic volatility and realized cumulative volatility, are very similar for high frequency data without structure noise. The results suggest that volatility at high frequency can be captured instantaneously for the medical measurements by using the Stochastic Volatility Model.

1. INTRODUCTION

Physiological measurements from preterm infants can display instabilities because of their under-developed biological systems. Besides considering the normal level of a physiological measurement, the variability of the measurement also supplies additional information about the underlying state of health. Estimation of the variability of certain physiological measurements from preterm infants is therefore of medical interest. When the medical preference is to acquire an instantaneous index of variability along with the measurement, the underlying variability has to be estimated with the aid of a suitable model. Although this seems to be uncommon in medical applications, it has been widely discussed by financial economists in the literature over the last two decades and referred to as volatility models.

The Stochastic Volatility and Autoregressive Conditional Heteroskedasticity class of models have become widely established and successful approaches to the modeling of the return variance process in financial time series. A variety of volatility models have been discussed by financial economists, such as the Autoregressive Conditional Heteroskedastic (ARCH) model of Engle (1982), the generalized ARCH (GARCH) model of Bollerslev (1986) and Taylor (1986), the exponential GARCH (EGARCH) model of Nelson (1991), and the stochastic volatility (SV) models of Melino and Turnbull (1990), and Harvey, Ruiz, and Shephard (1994). These models provide the latent volatility of the underlying financial product returns, as the volatility is not directly observable. With the move towards the use of high frequency data recently, there is interest in alternative volatility estimators, such as realized volatility (or integrated volatility). Particularly, French, Schwert and Stambaugh (1987) use daily returns to estimate monthly volatilities. Andersen and Bollerslev (1998) show that realized volatility computed from high-frequency intraday returns is effectively an error-free volatility measure. It is therefore natural to treat the consistent estimation of the realized volatility as observed.

This paper uses alternative volatility estimation approaches with an application on the oxygen concentration of preterm infants. Specifically, we apply the Stochastic Volatility Model (SVM) estimated by particle filters, the EGARCH model, and the method of Realized Volatility (RV) to get the observed volatilities as the approximation to the true volatility. We consider the observed integrated volatility as a benchmark to evaluate the performance of the latent volatility models.

2. ALTERNATIVE VOLATILITY ESTIMATIONS

2.1 SVM Model

In the paper, the stochastic volatility model is defined as the following:

$$\begin{aligned} z_t &= u + e^{x_t/2}\varepsilon_t \\ x_t &= c + ax_{t-1} + b\xi_t \end{aligned} \qquad (1)$$

where z_t is the measurement term and defined as $log(oxy_t/oxy_{t-1})$; The volatility x_t is defined as $log(\sigma^2)$ of the $log(oxy_t/oxy_{t-1})$ to ensure the positiveness of the conditional variance. Noise terms, ε_t and ξ_t, follow a standard normal distribution. The coefficient, *c* is a real number, *b* is a positive scale factor, and *a* is in the range (-1, 1).

Introducing the additional innovation term in the conditional mean process increases the flexibility of the model but it also increases the difficulty in parameter estimation because each shock is now described by two innovations, ε_t & ξ_t. Another difficulty with the SVM is the non-linear feature of the model. In this paper, the SVM is estimated using a particle filter algorithm, which is based on approximations in the representation of the desired distributions by discrete random measures. The standard assumptions are that the distributions of the two noise processes in the model are known, and continuous distributions are approximated by discrete random measures, which are composed of weighted particles of the unknown states (volatility x) and coefficients, with the weights computed using Bayes theory. Particle filters are particularly attractive for applications requiring on-line estimation.

2.2 EGARCH Model

Following the introduction of ARCH processes by Engle (1982) and their generalization by Bollerslev (1986), there have been numerous modifications of this approach to modeling conditional volatility. Nelson (1991) proposes the exponential GARCH (EGARCH) model, which allows for asymmetric effects between positive and negative values to overcome some weaknesses of the GARCH model in handling financial time series. We use the EGARCH model for comparison with SVM and RV. The following EGARCH process is used to model the relationship between the oxygen concentration changing measurements and their conditional volatilities:

$$
\begin{aligned}
z_t &= u + \varepsilon_t \\
\varepsilon_t &= e^{x_t/2} \delta_t \\
x_t &= a + b x_{t-1} + c\left[|\delta_{t-1}| - E(|\delta_{t-1}|)\right] + d(\delta_{t-1})
\end{aligned} \quad (2)
$$

As with the SVM, where z_t is the measurement term and defined as $log(oxy_t/oxy_{t-1})$ and ε_t serves as the shock from mean; The volatility x_t is defined as $log(\sigma^2)$ of the $log(oxy_t/oxy_{t-1})$; δ_t is standard Gaussian iid sequence and:

$$E(|\delta_t|) = E\left(\frac{|\varepsilon_t|}{e^{x_t/2}}\right) = \sqrt{2/\pi} \quad (3)$$

Maximum likelihood estimation is used to estimate the parameters.

2.3 Realized Volatility

As demonstrated by Andersen and Bollerslev (1998), Andersen (2001), and Barndorff-Nielsen and Shephard (2002), the sum of squared intraday returns is an approximation to the daily realized variance. This version of variation is named realized volatility (RV) and regarded as an estimation of integrated volatility. Unlike latent volatility as described above, RV is based on the information within a time window rather than all previous information until time t. RV can be built up easily as an integrated measurement of the z^2 at low frequency using high frequency intra-window data, and is a consistent estimator of observed volatility when high frequency data are available without structure bias and serial correlation. Forecasting and modeling volatility is equivalent to forecasting and modeling Realized Volatility.

However, this model-free estimation may not always be available as medical practitioners may require volatility estimates at a higher frequency than can be estimated reliably at the available data rate. In this paper, RV is used as a benchmark to evaluate the performance of alternative latent volatility models. At the same time, the empirical results can be used to evaluate the usefulness of this simple volatility method should large amounts of high frequency medical data be available in the future.

3. DATA

Three main types of time series are involved: raw oxygen data at every 2 second; estimated volatilities at every 2 seconds from the SVM and EGARCH model; observed volatility per minute from RV Method and cumulative latent volatility based on estimated volatilities over the same time window as RV under the two latent volatility models.

The raw oxygen dataset, comprising measurements every 2 seconds, was supplied by Christchurch Women's Hospital and was originally collected for an earlier study entitled, "Normal variation in oxygen levels in preterm babies". A segment of 4070 observations was taken for analysis. The variable z_t for $log(oxy_t/oxy_{t-1})$ is calculated based on the raw oxygen data and used by both latent volatility models and realized volatility. Using this dataset, the time varying conditional volatilities at 2-second intervals are estimated based on the discrete-time EGARCH and stochastic volatility models.

The RV is obtained for every minute using the 2-second data, giving 135 volatility estimates. We also looked at different time-frequency to check the robustness of the estimated realized volatility. Unlike financial data, the medical data do not have microstructure bias but still has slight serial correlation problems for intra-period data.

In order to compare time-varying volatility under SVM and EGARGH model with RV, cumulative volatility per minute is calculated based on the estimated volatility per 2 second for both latent volatility models. Table 1 gives the summary statistics of the results.

Table 1. Descriptive Statistics

	Zt(oxygen)	Volatility (per2second)	
		SVM	EGARCH
Maximum	0.2315	0.1663	0.1923
Minimum	-0.1671	0.0011	0.0031
Mean	0	0.0187	0.0190
Std.Dev	0.0234	0.0160	0.0134

Table 1. Continued

	Volatility (per minute)		
	SVM	EGARCH	RV
Maximum	0.3192	0.4148	0.3297
Minimum	0.0192	0.0285	0.0229
Mean	0.1142	0.1117	0.1071
Std.Dev	0.0722	0.0616	0.0709

4. PERFORMANCE OF ALTERNATIVE VOLATILITY MODELS

Both the SVM and EGARCH use $log(\sigma^2)$ instead of σ^2 to ensure positiveness of the conditional variance. Figure 1 shows the $log(\sigma^2)$ time series estimated under SVM using a particle filter. There are clear volatility clusters in the estimates. EGARCH estimation of $log(\sigma^2)$ has a similar pattern.

The time series plot of volatilities (σ) from both latent volatility models is shown in Figure 2. Although the two sequences of estimates appear to be consistent over the entire time series, there are clear differences particularly for low and high volatilities.

In order to compare the performance of the two latent volatility models, we use RV as a benchmark for comparison. Since RV is available only every minute, the 2-second estimates of latent volatility must be transformed to latent volatility per minute. The proportion of the total variation of RV that can be explained by the estimated latent volatility can then be measured using R^2 from a least-squares linear regression analysis:

$$RV_t = \alpha_0 + \alpha_1 SVM_{t/t-1} + \delta_t \quad (4)$$

$$RV_t = \beta_0 + \beta_1 EGARCH_{t/t-1} + \delta_t \quad (5)$$

The R^2 value for SVM volatility is about 93.95%, while for EGARCH, it is about 81.36%. Thus, the R^2 for SVM is about 13% higher than for EGARCH. This suggests that SVM volatility describes the dynamics of RV better than EGARCH volatility.

In addition to R^2, we compute the mean absolute error (MAE) and root mean squared error (RMSE), as well as the heteroscedasticity-adjusted mean absolute error (HMAE) and root mean squared error (HRMSE), to measure how far the estimated latent volatilities are from the realized volatility. These measures are defined as follows:

$$MAE = \frac{1}{n}\sum_{t=1}^{n}\left| vol_{estimated,t} - vol_{realized,t} \right| \quad (6)$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(vol_{estimated,t} - vol_{realized,t})^2} \quad (7)$$

$$HMAE = \frac{1}{n}\sum_{t=1}^{n}\left|(1 - \frac{vol_{realized,t}}{vol_{estimated,t}})\right| \quad (8)$$

$$HRMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(1 - \frac{vol_{realized,t}}{vol_{estimated,t}})^2} \quad (9)$$

These four performance measures are reported in Table 2. It is clear that both the MAE and RMSE of SVM are smaller than EGARCH. After the deviation between estimated latent volatility and realized volatility is adjusted for heteroscedasticity, SVM still outperforms EGARCH with smaller HMAE and HRMSE.

Table 2. Performance of alternative volatility models

	MAE	RMSE	HMAE	HRMSE
EGARCH	0.0206	0.0309	0.1829	0.2250
SVM	0.0123	0.0191	0.1207	0.1546

Further useful insight can be gain from Figure 3, which presents volatility per minute for SVM, EGARCH and RV together. SVM volatility is closer to RV than compared to EGARCH volatility.

We further obtain approximate 95% confidence intervals for RV at every minute by a bootstrap method, to assess how close SVM volatility estimates are to RV. The results are shown in Figure 4. The dashed line with circles is the confidence interval for each time-point. The star symbol represents realized volatility every minute and the dashed line with plus sign is the SVM volatility per minute. There are few points outside the confidence intervals, and these occur within the first ten time-points, corresponding to the initial period where the particle filter has not yet converged. The estimates from the particle filter improve as more measurements are processed, and so the estimates are closer together as time progresses. HMAE and HRMSE are 0.1288 and 0.1577 for the first half of the data and smaller HMAE and HRMSE are

0.1062 and 0.1363 for the second half of the data.

5. CONCLUSION

This paper models the conditional volatility of oxygen concentration of preterm infants using financial volatility models – SVM and realized volatility. The empirical results suggest that volatility estimation using SVM is consistent with observed volatility, and therefore may be used to estimate the instantaneous latent volatility for the oxygen concentration data.

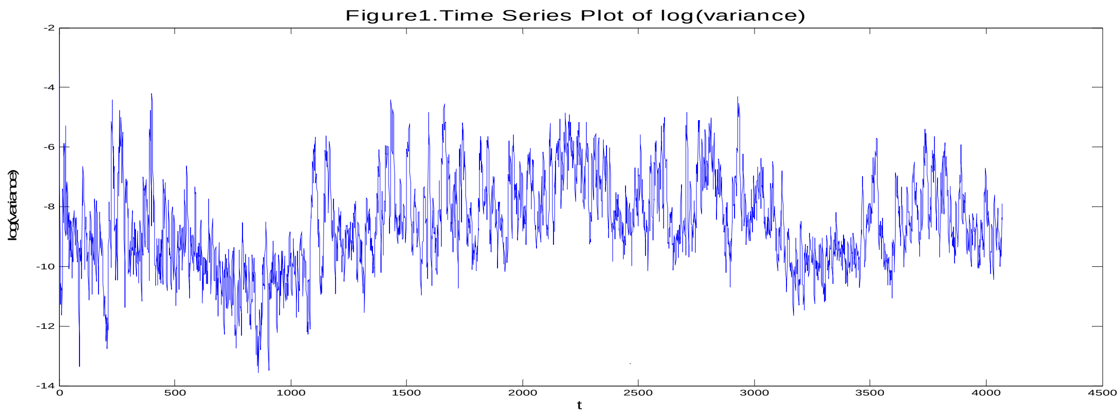

Figure 1. Time Series Plot of log(σ^2) estimates for oxygen measurements based on SVM model with particle filter. Length of series is 4069, based on 4070 oxygen measurements at 2-second intervals.

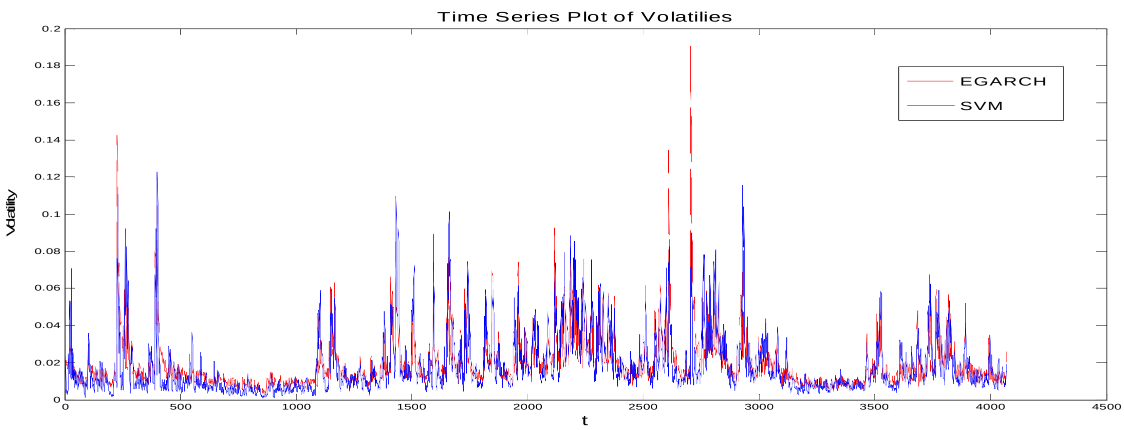

Figure 2. Time Series plot of estimated volatility (σ) based on SVM and EGARCH model.

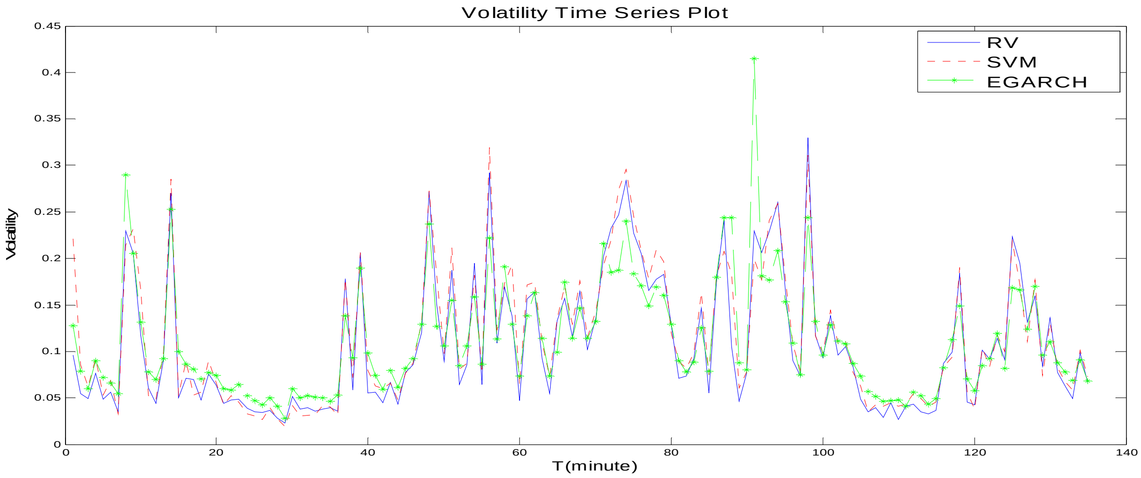

Figure 3. Time series plot of latent volatility by SVM and EGARCH with observed volatility RV.

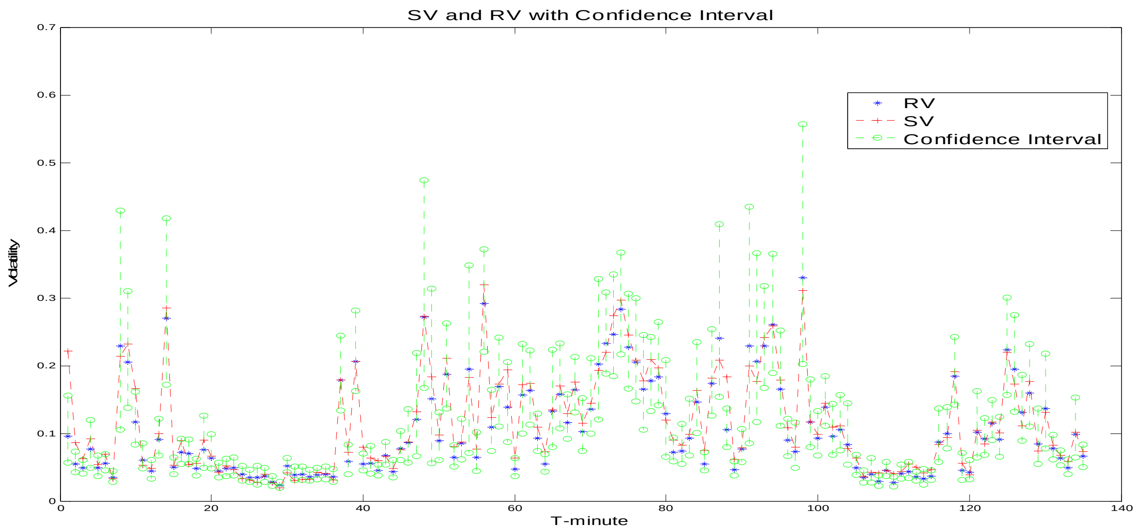

Figure 4. SVM volatility and Realized volatility with confidence intervals by bootstrap.

6. REFERENCES

Andersen, T.G., T. Bollerslev (1998), Deutsche Mark-Dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance,* 6, 457-477.

Andersen, T.G., T. Bollerslev, F.X. Diebold, and H. Ebens (2001), The distribution of realized stock return volatility, *Journal of Financial Economics,* 61, 43-76.

Andersen, T.G., T. Bollerslev and N. Meddahi (2005), Correcting the errors: volatility forecast evaluation using high-frequency data and realized volatilities. *Econometrica,* 73, 279-296.

Andreou, E. and E. Ghysels (2002), Rolling-sample volatility estimators: some new theoretical, simulation and empirical results, *Journal of Business and Economic Statistics,* 20, 363-376.

Bali, T.G. and D. Weinbaum (2007), A conditional extreme value volatility estimator based on high-frequency returns, *Journal of Economic Dynamics & Control*, 31, 361-397

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

Barndorff-Nielsen, O.E. and N. Shephard, (2002), Econometric analysis of realized volatility and its use in estimating stochastic volatility models, *Journal of the Royal Statistical Society*, B64, 253-280.

Beresford, M.W., H. Parry and N.J. Shaw (2005), Twelve-month prospective study of oxygen saturation measurements among term and preterm infants, *Journal of Perinatology*, 25, 30-32.

Chen, T., J. Morris and E. Martin, (2005), Particle filters for state and parameter estimation in batch processes, *Journal of process control*, 15, 665-673.

Djurić, P.M., J.H. Kotecha, J. Zhang, Y. Huang, T. Gbirmai, M.F. Bugallo and J. Míguez (2003), *IEEE signal processing magazine*, September, 19-38.

Fleming, J., B. Ostdiek and R.E. Whaley (1995), Predicting stock market volatility: a new measure, *Journal of Futures Markets*, 15, 256-302.

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

Frances, P.H. and D. Van Dijk (1995), Forecasting stock market volatility using (non-linear) GARCH models, *Journal of Forecasting*, 15, 229-235.

Hentschel, L.E., (1995), All in the family: nesting symmetric and asymmetric GARCH models, *Journal of Financial Economics*, 39, 71-104.

Magni, P. and R. Bellazzi (2006), A Stochastic Model to assess the variability of blood glucose time series in diabetic patients self-monitoring, *IEEE Transactions on Biomedical Engineering*, Vol 53, No 6, 977-985.

McAller, M. and M.C. Medeiros (2006), Realized volatility: A review, *Upublished manuscript, School of Economics and Commerce, University of Western Australia and Department of Economics, Pontifical Catholic University of Rio de Janeiro*

Nelson, D.B. (1991), Conditional heteroscedasticity in asset returns: a new approach, *Econometrica*, 59, 347-370.

Ng, A., N. Subhedar, R.A. Primhak and N.j. Shaw (1998), Arterial oxygen saturation profiles in healthy preterm infants, *Arch. Dis. Child. Fetal Neonatal Ed*, 79, 64-66.

Oomen, R.C.A. (2002), Modelling Realized Variance when Returns are serially correlated. *Upublished manuscript, Graduate School of Business School, University of Warwick*.

Sanjeev Arulampalam, M., S. Maskell, N. Gordon and T. Clapp (2002), A tutorial on particle filters for online Nonlinear/Non-Gaussian Bayesian tracking, *IEEE Transactions on Signal Processing*, 50, No.2, February, 174-188.

Tsay, R.S. (2002), Analysis of Financial Time Series.