

NCER Working Paper Series

A nonparametric approach to forecasting realized volatility

Adam Clements
Ralf Becker

Working Paper #43
May 2009

A nonparametric approach to forecasting realized volatility

A Clements

School of Economics and Finance, Queensland University of Technology

R Becker

Economic Studies, University of Manchester.

Abstract

A well developed literature exists in relation to modeling and forecasting asset return volatility. Much of this relate to the development of time series models of volatility. This paper proposes an alternative method for forecasting volatility that does not involve such a model. Under this approach a forecast is a weighted average of historical volatility. The greatest weight is given to periods that exhibit the most similar market conditions to the time at which the forecast is being formed. Weighting occurs by comparing short-term trends in volatility across time (as a measure of market conditions) by the application of a multivariate kernel scheme. It is found that at a 1 day forecast horizon, the proposed method produces forecasts that are significantly more accurate than competing approaches.

Keywords

Volatility, forecasts, forecast evaluation, model confidence set, nonparametric

JEL Classification Numbers

C22, G00.

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 the volatility of asset returns are of great interest to many financial market participants. Applications such as risk management, portfolio allocation and derivative pricing all require such forecasts. There has been a vast literature relating to forecasting asset return volatility. Much of this has focused on the development of econometric models of volatility, surveys of which can be found in Campbell, Lo and MacKinlay (1997) and Gouriéroux and Jasiak (2001). Much of this literature has stemmed from the development of the GARCH class of models attributable to Engle (1982) and Bollerslev (1986). Recently, it has benefited from the development of Realized Volatility (RV) by Andersen, Bollerslev, Diebold and Labys (2001, 2003). Time series models such as the Mixed Interval Data Sampling (MIDAS) framework have been directly applied to RV estimates for forecasting purposes. An alternative method for obtaining forecasts is to rely on implied volatility (IV), derived from option prices. IV should represent a market's best prediction of an assets' future volatility (see, amongst others, Jorion, 1995, Poon and Granger, 2003, 2005). Poon and Granger (2003, 2005) provide a wide ranging survey of articles relating to forecasting volatility. While the results are somewhat mixed, overall, option based forecasts are often more accurate than those based on econometric models.

This paper proposes a nonparametric approach to forecasting realized volatility. The principle is similar in nature to that of Brandt (1999) however the mechanics in this context are quite different. The proposed approach begins by measuring short-term trends in volatility as a measure of market conditions. A forecast is then given by a weighted average of historical RV, where the greatest weight is given to periods that are most similar to the time at which the forecast is being formed. Weights are obtained by the application of a multivariate kernel, while historical observations of both RV and IV are used to capture market conditions. While results pertaining to IV as a forecast in its own right are mixed, it has been found to be a useful measure of market volatility in conjunction with RV. The performance of the proposed kernel based forecast will be compared to IV forecasts, a model based solely on daily returns, and a number of time series models that utilize RV. It is found that at a 1 day horizon, the kernel based forecast is statistically superior to the competing models. At longer horizons, the performance of the proposed approach is equivalent to a number of alternative models.

The paper proceeds as follows. Section 2 describes the data used in this study. Specifically, the daily returns data, intraday data upon which RV is based and the IV estimates. Section 3 will outline the proposed kernel based forecast, competing forecasts and the manner in which their performance will be evaluated. Sections 4 and 5 report the empirical results and provide concluding remarks respectively.

2 Data

This study utilizes data relating to the S&P 500 Composite Index, from 2 January 1990 to 31 October 2008 equating to 4791 daily observations. Daily index return data, IV and RV estimates are required for the current analysis.

The *VIX* index constructed by the Chicago Board of Options Exchange from S&P 500 index options constitutes the estimate of IV utilized in this paper. It is derived from out-of-the-money put and call options that have maturities close to the target of 22 trading days¹. The *VIX* is constructed to be a general measure of the market's estimate of average S&P 500 volatility over the subsequent 22 trading days BPT, 2001, Christensen and Prabhala, 1998 and CBOE, 2003. Having a fixed forecast horizon is advantageous and avoids various econometric issues. This index has only been available since September 2003 when the CBOE replaced a previous implied volatility index based on S&P 100 options². Its advantages in comparison to the previous implied volatility index is that it no longer relies on option implied volatilities derived from Black-Scholes option pricing models, it is based on more liquid options written on the S&P500 and is easier to hedge against (CBOE, 2003).

For the purposes of this study estimates of actual volatility were obtained using the RV methodology outlined in ABDL (2001, 2003). 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 returns³. It is widely acknowledged that RV is a 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 shows the *VIX* and daily S&P500 RV for the sample period considered. While RV estimates exhibit a similar overall pattern when compared to the *VIX*, RV reaches higher peaks than the *VIX*. This difference is mainly due to the fact that the *VIX* represents an average volatility measure for a 22 trading day period as opposed to RV that is a measure of daily volatility. The high volatility period of September to December 2008 is clearly evident.

¹For technical details relating to the construction of the *VIX* index, see Chicago Board Options Exchange CBOE, 2003.

²The new version of the *VIX* has been calculated retrospectively back to January 1990, the beginning of the sample considered here.

³Intraday S&P 500 index data were purchased from Tick Data, Inc.




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

During this time, both RV and the VIX reached unprecedented levels as equity markets fell as a consequence of the credit crisis and ensuing financial turmoil.

3 Methodology

This section will begin by outlining the details of the proposed nonparametric forecast. This will be followed by a brief description of the forecasts with which it will be compared, along with the technique utilized in evaluating the forecasts.

3.1 A nonparametric forecast

As briefly discussed in Section 1, the nonparametric forecast is based on a weighted average of historical RV. The greatest weight is given to periods that are most similar in terms of market conditions to the time at which the forecast is being formed. Market conditions at time τ are captured by short-term moving averages in historical RV and IV which we shall define as

$$\Phi_\tau = [\overline{RV}_\tau^{(\lambda_1)}, \dots, \overline{RV}_\tau^{(\lambda_p)}, \overline{VIX}_\tau^{(\lambda_1)}, \dots, \overline{VIX}_\tau^{(\lambda_p)}]', \quad (1)$$

where $\overline{RV}_\tau^{(\lambda_i)}$ and $\overline{VIX}_\tau^{(\lambda_i)}$ are λ_i period moving averages (ending) at time τ and values for λ_i have been selected to be $\lambda_i = 1, 2, 3, 4, 5, 7, 10$.

Assume that at time t we are to forecast volatility over the subsequent period of q days (i.e. over days $t+1$ to $t+q-1$). The forecast will be a weighted average of all available $\overline{RV}_\tau^{(q)}$, $\tau \leq t$. Collect all $\overline{RV}_\tau^{(q)}$, $\tau \leq t$, in a vector $\overline{\mathbf{RV}}_t^{(q)}$, such that the forecast at time t for the subsequent period of q days is $\mathbf{w}'\overline{\mathbf{RV}}_t^{(q)}$.

The weights associated with each $\overline{RV}_\tau^{(q)}$ will be determined by the similarity of the market conditions on the day before the start of the q period covered by $\overline{RV}_\tau^{(q)}$, $\Phi_{\tau-q}$, with those pertaining at time t , Φ_t . Given $\Phi_{\tau-q}$, the weight attached to $\overline{RV}_\tau^{(q)}$ is given by a multivariate product kernel,

$$\tilde{w}_{\tau-q} = \prod_{n=1}^N K \left\{ \frac{\Phi_{t,n} - \Phi_{\tau-q,n}}{h_n} \right\}, \quad (2)$$

where K is a standard normal kernel, $\Phi_{t,n}$ is the n th element in Φ_t and N is the number of dimensions in Φ . Based on the optimal bandwidths for multivariate density estimation of Scott (1992), the bandwidth for dimension n , h_n is given by

$$h_n = \sigma_n T^{\frac{-1}{4+N}} \quad (3)$$

where σ_n is the standard deviation of the elements of dimension n in Φ and T is the corresponding number of observations of Φ available in the estimation period⁴. The weight vector \mathbf{w} is then scaled

$$\mathbf{w} = \frac{\tilde{\mathbf{w}}_{\tau-q}}{\tilde{\mathbf{w}}_{\tau-q}'\boldsymbol{\iota}} \quad (4)$$

where $\boldsymbol{\iota}$ is a vector of ones, ensuring that the elements in \mathbf{w} sum to 1.

This nonparametric weighting mechanism does not, per se, require any parameter estimation, although, as with any nonparametric procedure, it is necessary to choose the smoothing parameter h_n .

3.2 Competing forecasts

The kernel based forecast will be compared to a number of alternatives, including IV, in the form of the VIX along with a number of model based forecasts. The simplest begins with the GJR model of Glosten, Jagannathan and Runkle (1993) based on daily return observations,

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 \varepsilon_{t-1}^2 I_{\varepsilon_{t-1} < 0} + \alpha_3 \sigma_{t-1}^2 \quad (5)$$

where ε_t is a residual from a conditional mean equation and $I_{\varepsilon_t < 0}$ is an indicator taking the value of 1 if $\varepsilon_t < 0$. This simple model is also extended to include RV_{t-1} as an exogenous regressor in equation 5 and will be denoted below as GJR^{RV} . The final forecast considered

⁴As we are not estimating a density function, the kernel weighting scheme is a simple data dependent approach for obtaining weights

here is a MIDAS forecast, given a direct time-series model of RV, see Ghysels, Santa-Clara and Valkanov (2006). A MIDAS forecast is generated according to⁵

$$\overline{RV}_{t+1 \rightarrow t+q} = \sum_{k=0}^{k_{\max}} b(k, \theta) RV_{t-k} + \varepsilon_t \quad (6)$$

The maximum lag length k_{\max} can be chosen rather liberally as the weight parameters $b(k, \theta)$ are tightly parameterized. In this case $k_{\max} = 200$ is chosen. Here the weights are determined by means of a beta density function and normalized such that $\sum b(k, \theta) = 1$. A beta distribution function is fully specified by the 2×1 parameter vector θ . Parameter estimation was achieved by nonlinear least squares, minimizing the sum of squared residuals in equation (6).

3.3 Evaluating volatility forecasts

The Model Confidence Set approach (MCS) of Hansen, Lunde and Nason (2003) will be used to evaluate the forecast performance of the competing models. The MCS is a modified version of the Superior Predictive Ability (SPA) test of Hansen (2005) in that it has greater power and does not require a benchmark forecast to be chosen. Application of the MCS produces a set of models that are statistically indistinguishable in terms of their forecast performance.

The procedure starts with a full set of candidate models $\mathcal{M}_0 = \{1, \dots, m_0\}$. The MCS is determined by sequentially trimming models from \mathcal{M}_0 therefore reducing the number of models to $m < m_0$. Prior to starting the sequential elimination procedure, all loss differentials between forecasts i and j , $h_{t+1 \rightarrow t+q}^i$ and $h_{t+1 \rightarrow t+q}^j$ are computed,

$$d_{ij,t+1 \rightarrow t+q} = L(\hat{\sigma}_{t+1 \rightarrow t+q}^2, h_{t+1 \rightarrow t+q}^i) - L(\hat{\sigma}_{t+1 \rightarrow t+q}^2, h_{t+1 \rightarrow t+q}^j), \quad (7)$$

$\forall i, j = 1, \dots, m_0$ and $t = T_1, \dots, T_2 - q$. In this case, T_1 represents the final observation prior to the first forecast period and T_2 is the final observation in the dataset. The volatility proxy, $\hat{\sigma}_{t+1 \rightarrow t+q}^2$ is $\overline{RV}_{t+1 \rightarrow t+q}$ and the forecasts are those described in Sections 3.1 and 3.2. Mean Squared Error (MSE) and the Quasi-Likelihood (QLIKE) loss functions, $L(\cdot, \cdot)$ in equation 7 will be used within the MCS,

$$\begin{aligned} MSE &= (\hat{\sigma}_{t+1 \rightarrow t+q}^2 - h_{t+1 \rightarrow t+q})^2 \\ QLIKE &= \log(h_{t+1 \rightarrow t+q}) + \frac{\hat{\sigma}_{t+1 \rightarrow t+q}^2}{h_{t+1 \rightarrow t+q}}. \end{aligned} \quad (8)$$

Patton (2006) proved that while many loss functions exist, MSE and QLIKE are commonly used loss functions that belong to a family of loss functions are robust to noise in the volatility proxy.

⁵MIDAS 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 directly. These generalizations are not required here.

At each step, the EPA hypothesis

$$H_0 : E(d_{ij,t+1 \rightarrow t+q}) = 0, \quad \forall i > j \in \mathcal{M} \quad (9)$$

is tested for a set of models $\mathcal{M} \subset \mathcal{M}_0$, with $\mathcal{M} = \mathcal{M}_0$ at the initial step. If H_0 is rejected at the significance level α , the worst performing model is removed and the process continued until non-rejection occurs with the set of surviving models being the MCS, $\widehat{\mathcal{M}}_\alpha^*$. If a fixed significance level α is used at each step, $\widehat{\mathcal{M}}_\alpha^*$ contains the best model from \mathcal{M}_0 with $(1 - \alpha)$ confidence⁶.

At the core of the EPA statistic is the t -statistic

$$t_{ij} = \frac{\bar{d}_{ij}}{\sqrt{\widehat{\text{var}}(\bar{d}_{ij})}}, \quad (10)$$

where $\bar{d}_{ij} = \frac{1}{T_2 - q - T_1 + 1} \sum_{t=T_1}^{T_2 - q} d_{ij,t+1 \rightarrow t+q}$. t_{ij} provides scaled information on the average difference in the forecast quality of models i and j . $\widehat{\text{var}}(\bar{d}_{ij})$ is an estimate of $\text{var}(\bar{d}_{ij})$ and is obtained from a bootstrap procedure⁷. In order to decide whether, at any stage, the MCS must be further reduced, the null hypothesis in equation 9 is to be evaluated. The difficulty being that for each set \mathcal{M} the information from $(m - 1) m/2$ unique t -statistics needs to be distilled into one test statistic. Hansen *et al.* (2003) propose the following the range statistic,

$$T_R = \max_{i,j \in \mathcal{M}} |t_{ij}| = \max_{i,j \in \mathcal{M}} \frac{|\bar{d}_{ij}|}{\sqrt{\widehat{\text{var}}(\bar{d}_{ij})}} \quad (11)$$

and a semi-quadratic statistic,

$$T_{SQ} = \sum_{\substack{i,j \in \mathcal{M} \\ i < j}} t_{ij}^2 = \sum_{\substack{i,j \in \mathcal{M} \\ i < j}} \frac{(\bar{d}_{ij})^2}{\widehat{\text{var}}(\bar{d}_{ij})} \quad (12)$$

as test statistics to establish EPA. Both test statistics indicate a rejection of the EPA hypothesis for large values. The actual distribution of the test statistic is complicated and depends on the covariance structure between the forecasts included in \mathcal{M} . Therefore p-values for each of these test statistics have to be obtained from the bootstrap distribution. When the null hypothesis of EPA is rejected, the worst performing model is removed from \mathcal{M} . The latter is identified as \mathcal{M}_i where

$$i = \arg \max_{i \in \mathcal{M}} \frac{\bar{d}_i}{\sqrt{\widehat{\text{var}}(\bar{d}_i)}} \quad (13)$$

and $\bar{d}_i = \frac{1}{m-1} \sum_{j \in \mathcal{M}} \bar{d}_{ij}$. The tests for EPA are then conducted on the reduced set of models and one continues to iterate until the null hypothesis of EPA is not rejected. Thus, the

⁶Despite the testing procedure involving multiple hypothesis tests this interpretation is a statistically correct one. See Hansen *et al.* (2003) for a detailed discussion of these aspects.

⁷For specific details on the bootstrap procedure see Becker and Clements (2008) and Hansen *et al.* (2003)

	<i>MSE</i>		<i>QLIKE</i>		
	T_R	T_{SQ}	T_R	T_{SQ}	
<i>GJR</i>	0.4750	0.4880	<i>VIX</i>	0.0000	0.0000
<i>Kern</i>	0.6130	0.5960	<i>GJR</i>	0.0000	0.0000
<i>GJR^{RV}</i>	0.8010	0.7420	<i>MID^{RV}</i>	0.0100	0.0120
<i>MID^{RV}</i>	0.8010	0.7420	<i>GJR^{RV}</i>	0.0220	0.0220
<i>VIX</i>	1.0000	1.0000	<i>Kern</i>	1.0000	1.0000

Table 1: MCS results for 1 day forecasts of volatility. p-values given both the T_R and T_{SQ} test statistics are reported for both MSE and QLIKE loss functions.

final set of models constituting the MCS are models whose forecast performance is statistically indistinguishable.

4 Empirical Results

The forecasts will be evaluated at horizons of 1, 5 and 22 trading days. All models, including the kernel forecast were initially estimated on the first 1000 observations. A recursive estimation scheme was implemented with the estimation window extended by one day leading to 3791, 3787 and 3770, 1, 5 and 22 day ahead forecasts respectively. MCS results will be presented for each of the forecast horizons and are contained in Tables 1, 2 and 3 for the 1,5 and 22 day horizons respectively.

Results in Table 1 indicate that based on the MSE loss function, all of the models are statistically indistinguishable given the relatively high p-values. Given the QLIKE loss function, the result is quite different, the proposed kernel based forecast is the sole model in MCS (*GJR^{RV}* is rejected from the MCS at a p-value of 0.0220). Patton and Sheppard (2006) show that QLIKE, relative to MSE exhibits significantly more power in differentiating between forecasts. Given the MCS results of Table 1, it appears as though the kernel based method generates significantly superior forecasts at the 1 day horizon as it is the sole model in the MCS under QLIKE. Results are similar for the 5 day forecast horizon, as reported in Table 2. MSE cannot distinguish between any of the forecasts whereas under QLIKE the MCS contains 3 forecasts, one of which is the proposed kernel method. At the 22 day however, there is little difference between the performance of all of the forecasts. Results in Table 3 show that once again MSE cannot discriminate between any of the forecasts, and under QLIKE only the VIX forecast is identified as inferior. Thus overall, at very short forecast horizons, the proposed nonparametric approach provides forecast performance gains relative to a number of common alternatives.

	<i>MSE</i>		<i>QLIKE</i>		
	T_R	T_{SQ}	T_R	T_{SQ}	
<i>Kern</i>	0.4880	0.5180	<i>VIX</i>	0.0000	0.0000
<i>GJR</i>	0.6320	0.7200	<i>GJR</i>	0.0360	0.0550
<i>MID^{RV}</i>	0.7740	0.7680	<i>GJR^{RV}</i>	0.7710	0.8000
<i>VIX</i>	0.7740	0.7680	<i>MID^{RV}</i>	0.9640	0.9640
<i>GJR^{RV}</i>	1.0000	1.0000	<i>Kern</i>	1.0000	1.0000

Table 2: MCS results for 5 day forecasts of volatility. p-values given both the T_R and T_{SQ} test statistics are reported for both MSE and QLIKE loss functions.

	<i>MSE</i>		<i>QLIKE</i>		
	T_R	T_{SQ}	T_R	T_{SQ}	
<i>MID^{RV}</i>	0.2940	0.2670	<i>VIX</i>	0.0600	0.1630
<i>Kern</i>	0.2940	0.2670	<i>Kern</i>	0.8640	0.8560
<i>VIX</i>	0.2940	0.2670	<i>MID^{RV}</i>	0.8640	0.8560
<i>GJR</i>	0.2940	0.2670	<i>GJR^{RV}</i>	0.8640	0.8560
<i>GJR^{RV}</i>	1.0000	1.0000	<i>GJR</i>	1.0000	1.0000

Table 3: MCS results for 22 day forecasts of volatility. p-values given both the T_R and T_{SQ} test statistics are reported for both MSE and QLIKE loss functions.

5 Conclusion

This paper proposed a novel nonparametric technique for forecasting volatility. The forecast is a weighted average of historical RV, where the greatest weight is given to periods that exhibit the most similar market conditions to the time at which the forecast is being formed. Weighting occurs by comparing short-term trends in volatility across time (as a measure of market conditions) by the application of a multivariate kernel scheme.

It has been found that by utilizing historical RV and IV in determining market conditions, the proposed approach can produce significantly superior forecasts at a 1 day horizon. While Becker and Clements (2008) find that the VIX index is an inferior forecast in its own right, it seems to contain useful information about the state of market volatility.

The short-term forecast performance of the kernel approach may be attributable to the fact that the forecast is not simply a smoothed function historical RV. By taking a weighted average of RV we are not smoothing out potentially useful information such as the jump component of total volatility.

References

- 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.
- Becker, R., and Clements, A. (2008), "Are combination forecasts of S&P 500 volatility statistically superior?", *The International Journal of Forecasting*, 24, 122-133.
- Blair B.J., Poon S-H. & Taylor S.J. (2001). Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics*, 105, 5-26.
- Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, 31, 307-327.
- Brandt, M.W. (1999), Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach, *Journal of Finance*, 54, 1609-1646.
- Campbell, J.Y., Lo, A.W. & MacKinlay, A.G. (1997). *The Econometrics of Financial Markets*, Princeton University Press, Princeton NJ.
- Becker, R. and Clements, A., (2008), Are combination forecasts of S&P 500 volatility statistically superior?, *The International Journal of Forecasting*, 24, 122-133.
- Blair, B.J., Poon, S-H., Taylor, S.J., 2001. Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics*, 105, 5-26.
- Chicago Board of Options Exchange, 2003. VIX, CBOE Volatility Index.
- Christensen, B.J., Prabhala, N.R., 1998. The relation between implied and realized volatility. *Journal of Financial Economics*, 50, 125-150.
- Engle, R.F. (1982), "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation", *Econometrica*, 50, 987-1007.

- 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.
- Glosten, L.R., Jagannathan, R and Runkle, D.E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance*, 48, 1779-1801.
- Gourieroux C. and Jasiak J. (2001). *Financial Econometrics*. Princeton University Press: Princeton NJ.
- Hansen, P. (2005), "A test for superior predictive ability", *Journal of Business and Economic Statistics*, **23**, 365-380.
- Hansen, P., Lunde, A. and Nason, J. (2003), "Choosing the best volatility models: the model confidence set approach", *Oxford Bulletin of Economics and Statistics*, **65**, 839-861.
- Jorion P. (1995). Predicting volatility in the foreign exchange market. *Journal of Finance*, 50, 507-528.
- Patton, A. (2006), Volatility Forecast Comparison Using Imperfect Volatility Proxies, Quantitative Finance Research Centre, University of Technology, Sydney, research Paper Series, n. 175.
- Patton, A. and Sheppard, K. (2007),. Evaluating Volatility Forecasts, in *Handbook of Financial Time Series*, Andersen, T.G., Davis, R.A., Kreiss, J.P. and Mikosch, T. eds., Springer-Verlag.
- Poon S-H. & Granger C.W.J. (2003). Forecasting volatility in financial markets: a review. *Journal of Economic Literature*, 41, 478-539.
- Poon S-H. & Granger C.W.J. (2005). Practical Issues in forecasting volatility. *Financial Analysts Journal*, 61, 45-56.
- Scott, D.W. (1992) *Multivariate Density Estimation: Theory, Practice and Visualization* John Wiley: New York.

List of NCER Working Papers

No. 42 ([Download full text](#))

Uwe Dulleck, Rudolf Kerschbamer and Matthias Sutter

The Economics of Credence Goods: On the Role of Liability, Verifiability, Reputation and Competition

No. 41 ([Download full text](#))

Adam Clements, Mark Doolan, Stan Hurn and Ralf Becker

On the efficacy of techniques for evaluating multivariate volatility forecasts

No. 40 ([Download full text](#))

Lawrence M. Kahn

The Economics of Discrimination: Evidence from Basketball

No. 39 ([Download full text](#))

Don Harding and Adrian Pagan

An Econometric Analysis of Some Models for Constructed Binary Time Series

No. 38 ([Download full text](#))

Richard Dennis

Timeless Perspective Policymaking: When is Discretion Superior?

No. 37 ([Download full text](#))

Paul Frijters, Amy Y.C. Liu and Xin Meng

Are optimistic expectations keeping the Chinese happy?

No. 36 ([Download full text](#))

Benno Torgler, Markus Schaffner, Bruno S. Frey, Sascha L. Schmidt and Uwe Dulleck

Inequality Aversion and Performance in and on the Field

No. 35 ([Download full text](#))

T M Christensen, A. S. Hurn and K A Lindsay

Discrete time-series models when counts are unobservable

No. 34 ([Download full text](#))

Adam Clements, A S Hurn and K A Lindsay

Developing analytical distributions for temperature indices for the purposes of pricing temperature-based weather derivatives

No. 33 ([Download full text](#))

Adam Clements, A S Hurn and K A Lindsay

Estimating the Payoffs of Temperature-based Weather Derivatives

No. 32 ([Download full text](#))

T M Christensen, A S Hurn and K A Lindsay

The Devil is in the Detail: Hints for Practical Optimisation

No. 31 ([Download full text](#))

Uwe Dulleck, Franz Hackl, Bernhard Weiss and Rudolf Winter-Ebmer
Buying Online: Sequential Decision Making by Shopbot Visitors

No. 30 ([Download full text](#))

Richard Dennis
Model Uncertainty and Monetary Policy

No. 29 ([Download full text](#))

Richard Dennis
The Frequency of Price Adjustment and New Keynesian Business Cycle Dynamics

No. 28 ([Download full text](#))

Paul Frijters and Aydogan Ulker
Robustness in Health Research: Do differences in health measures, techniques, and time frame matter?

No. 27 ([Download full text](#))

Paul Frijters, David W. Johnston, Manisha Shah and Michael A. Shields
Early Child Development and Maternal Labor Force Participation: Using Handedness as an Instrument

No. 26 ([Download full text](#))

Paul Frijters and Tony Beatton
The mystery of the U-shaped relationship between happiness and age.

No. 25 ([Download full text](#))

T M Christensen, A S Hurn and K A Lindsay
It never rains but it pours: Modelling the persistence of spikes in electricity prices

No. 24 ([Download full text](#))

Ralf Becker, Adam Clements and Andrew McClelland
The Jump component of S&P 500 volatility and the VIX index

No. 23 ([Download full text](#))

A. S. Hurn and V.Pavlov
Momentum in Australian Stock Returns: An Update

No. 22 ([Download full text](#))

Mardi Dungey, George Milunovich and Susan Thorp
Unobservable Shocks as Carriers of Contagion: A Dynamic Analysis Using Identified Structural GARCH

No. 21 ([Download full text](#)) (forthcoming)

Mardi Dungey and Adrian Pagan
Extending an SVAR Model of the Australian Economy

No. 20 ([Download full text](#))

Benno Torgler, Nemanja Antic and Uwe Dulleck
Mirror, Mirror on the Wall, who is the Happiest of Them All?

No. 19 ([Download full text](#))

Justina AV Fischer and Benno Torgler

Social Capital And Relative Income Concerns: Evidence From 26 Countries

No. 18 ([Download full text](#))

Ralf Becker and Adam Clements

Forecasting stock market volatility conditional on macroeconomic conditions.

No. 17 ([Download full text](#))

Ralf Becker and Adam Clements

Are combination forecasts of S&P 500 volatility statistically superior?

No. 16 ([Download full text](#))

Uwe Dulleck and Neil Foster

Imported Equipment, Human Capital and Economic Growth in Developing Countries

No. 15 ([Download full text](#))

Ralf Becker, Adam Clements and James Curchin

Does implied volatility reflect a wider information set than econometric forecasts?

No. 14 ([Download full text](#))

Renee Fry and Adrian Pagan

Some Issues in Using Sign Restrictions for Identifying Structural VARs

No. 13 ([Download full text](#))

Adrian Pagan

Weak Instruments: A Guide to the Literature

No. 12 ([Download full text](#))

Ronald G. Cummings, Jorge Martinez-Vazquez, Michael McKee and Benno Torgler

Effects of Tax Morale on Tax Compliance: Experimental and Survey Evidence

No. 11 ([Download full text](#))

Benno Torgler, Sascha L. Schmidt and Bruno S. Frey

The Power of Positional Concerns: A Panel Analysis

No. 10 ([Download full text](#))

Ralf Becker, Stan Hurn and Vlad Pavlov

Modelling Spikes in Electricity Prices

No. 9 ([Download full text](#))

A. Hurn, J. Jeisman and K. Lindsay

Teaching an Old Dog New Tricks: Improved Estimation of the Parameters of Stochastic Differential Equations by Numerical Solution of the Fokker-Planck Equation

No. 8 ([Download full text](#))

Stan Hurn and Ralf Becker

Testing for nonlinearity in mean in the presence of heteroskedasticity.

No. 7 ([Download full text](#)) (published)

Adrian Pagan and Hashem Pesaran

On Econometric Analysis of Structural Systems with Permanent and Transitory Shocks and Exogenous Variables.

No. 6 ([Download full text](#)) (published)

Martin Fukac and Adrian Pagan

Limited Information Estimation and Evaluation of DSGE Models.

No. 5 ([Download full text](#))

Andrew E. Clark, Paul Frijters and Michael A. Shields

Income and Happiness: Evidence, Explanations and Economic Implications.

No. 4 ([Download full text](#))

Louis J. Maccini and Adrian Pagan

Inventories, Fluctuations and Business Cycles.

No. 3 ([Download full text](#))

Adam Clements, Stan Hurn and Scott White

Estimating Stochastic Volatility Models Using a Discrete Non-linear Filter.

No. 2 ([Download full text](#))

Stan Hurn, J. Jeisman and K.A. Lindsay

Seeing the Wood for the Trees: A Critical Evaluation of Methods to Estimate the Parameters of Stochastic Differential Equations.

No. 1 ([Download full text](#))

Adrian Pagan and Don Harding

The Econometric Analysis of Constructed Binary Time Series.