Non-linear Filtering with State Dependant Transition Probabilities: A Threshold (Size Effect) SV Model

Clements, Adam[1] and **Scott White**

School of Economics and Finance
Queensland University of Technology, GPO Box 2434, Brisbane Qld 4001, Australia.
[1]Email: a.clements@qut.edu.au

Keywords: *Non-linear filtering, stochastic volatility, size effect, threshold*

EXTENDED ABSTRACT

Much research has focused upon the dynamics of the conditional volatility of financial asset returns. Broadly speaking there are two important features of the process underlying volatility. These may be described as either a sign effect, where the level of volatility is related to the sign of past returns or a size effect, where the dynamics of volatility are related to prevailing level of volatility.

A great deal of work, at least in the context of dealing with equity returns has examined the nature of the sign effect. This has become known as the leverage effect and has revealed that conditional volatility is higher subsequent to negative returns. The size effect on the other hand has received somewhat limited attention in that the dynamics of many models of volatility are independent of the level of volatility. What limited evidence has been documented, shows that the persistence of US equity market volatility falls as the level of shocks to returns increase.

This paper considers the size effect within an stochastic volatility framework, a common class of models for modeling conditional volatility. Estimation of the parameters of stochastic volatility models is relatively difficult due to the presence of a latent factor (conditional volatility). While many estimation techniques exist for simple stochastic volatility models, they may be categorised as either computationally simple and relatively inaccurate or computationally burdensome and relatively accurate. Given that this paper considers whether stochastic volatility dynamics are related to the level of volatility, an estimation scheme that is flexible but also computationally feasible must be utilised for estimation.

Non-linear filtering techniques have been applied to the estimation of stochastic volatility models. While this is a flexible framework, with few restrictive assumptions, it has not gained wide acceptance due to the associated computational costs. A non-linear filtering algorithm for the estimation of stochastic volatility models with size effects is proposed. Considering the computational cost issues, estimation here is utilises a computationally efficient non-linear filtering algorithm. This algorithm is based on a discretised non-linear filtering algorithm where the relevant continuous state-space (for unobserved volatility) is discretised and the latent variable treated as if it were a discrete-valued Markov process. While this discretised non-linear filtering algorithm has been developed for efficient estimation of standard stochastic volatility models, this paper proposes an extension that accommodates possible size effects.

To assess the significance of the size effect, an hypothesis testing framework is proposed. While a standard stochastic volatility model is nested within a model incorporating size effects, it is not possible to use standard likelihood ratio tests. This is due to the manner in which the state-space of the stochastic volatility model with size effects is partitioned using a threshold parameter. Under the standard stochastic volatility case, this threshold parameter is unidentified and thus standard likelihood ratio tests are not applicable. To overcome these problems, a bootstrap testing procedure is suggested. Under this approach, a likelihood ratio statistic (stochastic volatility with size relative to the standard stochastic volatility model) is replicated given data generated under the assumption of the standard stochastic volatility process. The likelihood ratio statistic is then computed given the actual datasets and then compared to the distribution of the statistics under the null hypothesis of the standard stochastic volatility model. Doing so, allows for an empirical p-value to test the significance of the size effect to be obtained.

Two empirical datasets are considered, daily returns on the S&P 500 composite and the Yen/USD exchange rate. In both the equity and currency return cases, volatility dynamics are found influenced by the level of prevailing volatility. When volatility is relatively low (high), volatility is extremely (not) persistent with little (a great deal of) noise. This size effect was found to be statistically significant in both instances.

1 INTRODUCTION

Modeling the distribution of financial asset returns is a critically important issue within areas such as risk management, portfolio construction and option pricing. To accurately capture the conditional distribution of returns it is necessary to capture time-variation of volatility.

Estimating the conditional distribution of asset returns can be attributed to the seminal work of Engle (1982) and Bollerslev (1987) who developed the ARCH and GARCH models respectively. A vast amount of literature exists in this field, summaries of which are contained in Bollerslev, Chou and Kroner (1992), Pagan (1996) and Campbell, Lo and MacKinlay (1997).

An alternative approach to GARCH style models are the Stochastic Volatility (SV) class of models which treat conditional volatility as a latent variable that follows its own stochastic process. While on a practical level, it is difficult to estimate the parameters of SV models, they are theoretically appealing. Clark (1973), Tauchen and Pitts (1983) and Andersen (1996) theoretically motivate SV models from the perspective of capturing stochastic changes in information flow.

There has been an enormous amount of attention paid to the specification of volatility dynamics. Broadly speaking, the important features of volatility dynamics may be due to either sign (relationship between volatility and the sign of past returns) and or size (relationship between volatility and the size of past returns or volatility) effects. Much work has been directed at dealing with the sign effect, specifically the asymmetric relationship between returns and volatility. Within the GARCH class of models, Nelson (1991), Hentschel (1995) and Ding, Engle and Granger (1993), among others, propose threshold style models to capture the sign effect. Within an SV framework, Harvey and Shephard (1996) and So, Lam and Li (2002) propose asymmetric SV models where volatility dynamics are dependant on the sign of past returns, once again to capture the leverage effect.

An alternative feature that has received much less attention is the size effect where volatility dynamics are dependent on the level of volatility. Friedman and Laibson (1989), Gouriéroux and Monfort (1992), Engle and Ng (1993) and Longin (1997), consider the size effect within GARCH style models. For instance, Friedman and Laibson (1989) find that the persistence in conditional volatility falls when shocks to US stock returns are large.

To the best of the authors knowledge, the size effect in volatility has not been considered within an SV framework. Partly, this may be due to the problems surrounding the estimation of such models. It is therefore the goal of this paper to develop a non-linear filtering algorithm that allows the dynamics of volatility (the latent variable) in an SV setting to be dependent on the current level of volatility. This non-linear filtering framework builds upon the non-linear filter approach to dealing with latent variable models proposed by Kitigawa (1987). Given the non-linear filtering framework, an hypothesis test is suggested to ascertain whether volatility dynamics are in fact dependant on the level of volatility.

This paper proceeds as follows. Section 2 introduces the concept of an SV model with size effects. Section 3 outlines the non-linear filtering estimation framework as it applies to a standard SV model, along with adjustments to capture the size effect. An approach to testing the significance of the size effect is also proposed. Section 4 presents an empirical application of the SV models with a size effect, showing this it is an important feature of the two time series considered. Section 5 provides concluding remarks.

2 STOCHASTIC VOLATILITY WITH SIZE EFFECTS

A stochastic volatility (SV) model considers that returns (the observed variable) $\{y_t\}_{t=1}^{T}$ are generated by,

$$y_t = \exp(x_t/2)\, u_t \qquad u_t \sim N\,(0,1) \tag{1}$$

where $x_t = \ln(\sigma_t^2)$. SV models treat x_t as an unobserved (latent) variable, following its own stochastic path, the simplest being an AR(1) process,

$$x_t = \alpha + \beta\, x_{t-1} + w_t \qquad w_t \sim N(0, \sigma_w^2) \tag{2}$$

where errors, u_t and w_t are assumed to be independent. These equations describe the standard SV model where volatility dynamics are independent of the current level of volatility.

To incorporate a size effect into the SV dynamics it is necessary to condition the volatility dynamics on the level of volatility. To do so, the state-space of x_t will be partitioned by the point τ into two adjoining regions each with their own distinct volatility dynamics. Two regions are selected in this context to reflect relatively high and low volatility. To allow for the size effect in SV dynamics equation 2 must be augmented

$$x_t = \alpha_s + \beta_s\, x_{t-1} + \sigma_{w,s} w_t, \tag{3}$$

where the subscript S denotes the index of the region containing x_{t-1}. If $x_{t-1} < \tau$ dynamics will be governed by $\theta_1 = (\alpha_1,\, \beta_1,\, \sigma_{w,1})$ otherwise if $x_{t-1} > \tau$ dynamics will be governed by $\theta_2 = (\alpha_2,\, \beta_2,\, \sigma_{w,2})$.

The following section will now outline the non-linear filtering framework employed to estimate to the standard SV model of equation 2 along with an extension to deal with the level dependence implied by equation 3.

3 SV ESTIMATION AND TESTING OF SIZE EFFECT

To estimate the parameters of the SV model incorporating the size effect, this paper builds upon the non-linear filtering framework pioneered in Kitigawa (1987). While many other approaches to estimating SV models exist (for summaries see Ghysels *et al.* 1996 and Shephard 1996) the non-linear filtering approach has been chosen in this setting as it provides the flexibility required when incorporating non-standard features such as the size effect.

Estimation of a latent variable process such as equation 2 within a non-linear filtering framework is based on a recursive, prediction-update algorithm. This approach requires two density functions to be defined and a number of integrals to be evaluated. Let $r\left(y_t|\, x_t, Y_{t-1}, \theta\right)$ be the conditional distribution of y_t on x_t (given equation 1), $q\left(x_t|\, x_{t-1}, Y_{t-1}, \theta\right)$ be the conditional distribution of x_t on x_{t-1} (given equation 2) and θ. The one-step ahead prediction of the distribution of x_t conditional on Y_{t-1}, $f\left(x_t\,|Y_{t-1}, \theta\right)$, is given by

$$f\left(x_t\,|Y_{t-1}, \theta\right) = \int_{-\infty}^{\infty} q(x_t\,|x_{t-1}, Y_{t-1}, \theta)\, f\left(x_{t-1}\,|Y_{t-1}, \theta\right) dx_{t-1}. \tag{4}$$

Once a new observation, y_t, is available, the probability distribution of the state variable at time t, conditional on information at time t, $f\left(x_t\,|Y_t, \theta\right)$, may now be obtained as

$$f\left(x_t|\, Y_t, \theta\right) = \frac{r(y_t|x_t, Y_{t-1}, \theta)\, f\left(x_t\,|Y_{t-1}, \theta\right)}{f\left(y_t|\, Y_{t-1}, \theta\right)}. \tag{5}$$

The denominator of equation (5) is the likelihood of observing y_t conditional on y_{t-1} and θ and may be computed as

$$f\left(y_t|\, Y_{t-1}, \theta\right) = \int_{-\infty}^{\infty} r(y_t|x_t, \theta)\, f\left(x_t|\, Y_{t-1}, \theta\right) dx_t \tag{6}$$

which may be optimised (for all observations) to permit maximum likelihood (ML) estimates of SV parameters to be obtained. The method to do so will now be discussed.

3.1 Estimation of SV models using DNF

The non-linear filtering algorithm utilised here solves equations 4 to 6 based on a discretisation of state-space. This algorithm is known as the discretised non-linear filter (DNF). This allows the likelihood function of a continuously valued latent variable process to be evaluated in a similar manner to Markov models for discrete valued time series, see MacDonald and Zucchini (1997). In doing so, this avoids the use of numerical integration or simulation schemes.

Under the DNF approach, the pdf of the latent variable, x, is approximated by computing the probability of observing x within a set of discrete intervals (a histogram) as opposed to the linear spline approach suggested by Kitagawa (1987). In discretising the state-space, N equal-width adjacent intervals are defined across $\overline{x} \pm 6\ \sigma_x$, bounded by $\{w^i\}_{i=1}^{N+1}$, and centered on the points $\{x^i\}_{i=1}^{N}$ where

$$x^i = \frac{w^i + w^{i+1}}{2}. \tag{7}$$

In general terms, the probability of observing x within the interval centered on x^i, i.e. $x \in (w^i, w^{i+1}]$ is given by

$$p(x \in (w^i, w^{i+1}]) = \int_{w^i}^{w^{i+1}} f\left(x\right) dx \approx p(x^i) \tag{8}$$

where $f\left(x\right)$ is the continuous probability distribution of the of the unobserved state variable x. The series of $\{p(x^i)\}_{i=1}^{N}$ represent a discretised approximation to the continuous distribution $f\left(x\right)$. Based on this discrete approximation, the DNF captures the evolution of the state variable through time given definitions of a time-invariant set of transition probabilities and a set of conditional likelihoods.

Transitional probabilities

Given that the state space is defined over N adjacent intervals it is possible to compute an $N \times N$ matrix of time-invariant transition probabilities, $\widehat{q}$. The elements of this matrix, $\widehat{q}^{\,i,j}\ \forall i, j = 1, ..., N$, represent the probability of x migrating from the interval centred on x^j at time $t-1$, to the interval centred on x^i at time t and is given by

$$\widehat{q}^{\,i,j} = \delta\, q\left(x_t^i|\, x_{t-1}^j, \theta\right) \tag{9}$$

where δ is the interval width. In the case where $q(.)$ is a normal distribution,

$$\widehat{q}^{\,i,j} = \frac{\delta}{\sqrt{2\pi\sigma_v^2}} \exp\left(\frac{-\left(x^i - \alpha - \beta x^j\right)^2}{2\sigma_v^2}\right). \tag{10}$$

Conditional likelihoods

The likelihood of observing y_t conditional on x being within each discrete interval is found. The $T \times N$ likelihood matrix containing elements, $\widehat{r}_t^i \ \forall i = 1, ..., N$, is defined by

$$\widehat{r}_t^i = r\left(y_t |\, x_t^i,\, \theta\right) \tag{11}$$

In the standard SV model of equation 2, $\widehat{r}_t^i$ is given by

$$\widehat{r}_t^i = \frac{1}{\sqrt{2\pi \exp(x^i)}} \exp\left(-\frac{y_t^2}{2\exp(x^i)}\right) \tag{12}$$

Based on this set of conditional likelihoods and the time-invariant matrix of transition probabilities, the DNF proceeds with the following steps.

Prediction Step

The predicted probability of observing $x \in (w^i, w^{i+1}]$ at time t is given by (a discrete approximation to 4):

$$\begin{aligned} P_t^i &= p(x_t^i |\, y_{t-1}, \theta) \\ &= \sum_{j=1}^{N} \widehat{q}^{\,i,j} \ \cdot U_{t-1}^j. \end{aligned} \tag{13}$$

Update Step

The updated probability of observing $x \in (w^i, w^{i+1}]$ at time t, is defined as (a discrete approximation to 5)

$$\begin{aligned} U_t^i &= p(x_t^i\, |y_t, \theta) \\ &= \frac{\widehat{r}_t^i \cdot P_t^i}{p(y_t |\, Y_{t-1},\theta)} \end{aligned} \tag{14}$$

Likelihood

The denominator of equation (14) is the likelihood of observing y_t, given by

$$p(y_t |\, Y_{t-1},\theta) = \sum_{i=1}^{N} \widehat{r}_t^i \cdot P_t^{i,j} \tag{15}$$

The log-likelihood used to generate ML estimates of θ are obtained directly from equation 15 and is given by

$$\ln L = \sum_{t=1}^{T} \ln[p(y_t |\, y_{t-1},\theta)]. \tag{16}$$

For the DNF to be initialised, the prediction of the state probabilities at time $t = 1$ need to be selected. The state probabilities are initialised by discretising the unconditional distribution of the state variable such that

$$P_1^i = \int_{w^i}^{w^{j+1}} f\left(x \,|\, \theta\right) dx \tag{17}$$

where

$$f\left(x |\, \theta\right) \sim N\left(\frac{\alpha}{(1-\beta)}, \frac{\sigma_w^2}{(1-\beta^2)}\right). \tag{18}$$

The manner in which this DNF framework can be augmented to incorporate size effects will now be discussed.

3.2 Estimation of SV models with a size effect using DNF

To capture a size effect in SV dynamics it is necessary to adjust both the nonlinear filtering framework of equation 4 through 6 and the the associated estimation procedure. In terms of the non-linear filtering equations, only the prediction equation, equation 4 must be adjusted to reflect the size effect,

$$\begin{aligned} &f\left(x_t\, |y_{t-1}, \theta\right) \\ &= \int_{-\infty}^{\infty} [I\ q(x_t\, |x_{t-1}, \theta_1) + (I-1)\ q(x_t\, |x_{t-1}, \theta_2)\,] \\ &\quad f\left(x_{t-1}\, |y_{t-1}, \theta_1,\theta_2\right) dx_{t-1} \end{aligned} \tag{19}$$

where $I = 1\ if\ x_{t-1} < \tau$, with SV dynamics being governed by $\theta_1 = (\alpha_1,\ \beta_1,\ \sigma_{w,1})$ otherwise $I = 0\ if\ x_{t-1} > \tau$ results in dynamics being governed by $\theta_2 = (\alpha_2,\ \beta_2,\ \sigma_{w,2})$. This specification is consistent with equation 3 in that the dynamics governing the evolution of volatility at any point in time is dependent on the current level of volatility.

To estimate an SV model with a size effect, three adjustments to the DNF are necessary. It is necessary to choose τ so as the state-space of x may be partitioned into two adjoining regions. Within each region, discrete intervals must be chosen so as to discretise state-space. Finally, it is necessary to adjust the transition probability matrix, $\widehat{q}^{\,i,j}$ to reflect the distinct volatility dynamics of each region.

Region Choice

Whilst the state space of x is theoretically infinite, as with the standard SV case it must be discretised into a finite number of intervals. As the size effect, requires the use of two state equations, it is not immediately obvious how to span the state space of x. However, it is reasonable to assume that after accounting for the size effect, volatility would lie within the same region. Thus it is proposed that the state-space be defined such that it spans the region implied by the ML estimates of the SV model. The first step is to find the ML estimates of the parameters of the SV model and compute

$$\begin{aligned} max(x) &= \frac{\widehat{\alpha}}{(1-\widehat{\beta})} + 6\frac{\widehat{\sigma}_w}{\sqrt{(1-\widehat{\beta}^2)}} \\ min(x) &= \frac{\widehat{\alpha}}{(1-\widehat{\beta})} - 6\frac{\widehat{\sigma}_w}{\sqrt{(1-\widehat{\beta}^2)}} \end{aligned} \tag{20}$$

The region defined between $min(x)$ and $max(x)$ is believed to span the relevant state-space.

To split the state space of x_t into two regions, the threshold point, τ is defined under a restriction that $\tau \in [\frac{\max(x)+11\min(x)}{12}, \frac{11\max(x)+\min(x)}{12}]$. This ensures that there is a non-trivial distance between τ and the edges of the discretisation, $min(x)$ and $max(x)$.

Interval Choice

It is now necessary to define the discretisation within each region, $[\min(x),\ \tau]$ and $[\tau, \max(x)]$. Define number of intervals in the upper and lower regions as

$$N_U = round\left(\frac{N\,(\max(x)-\tau)}{\max(x)-\min(x)}\right),\ N_L = N - N_U \tag{21}$$

with the interval widths in each region $\delta_U = (\max(x) - \tau)/N_U$ and $\delta_L = (\tau - min(x))/N_L$ respectively. Define a set of interval edgepoints to discretise state-space. These edgepoints are defined by $\{w^i\}_{i=1}^{N+1} = \min(x), \min(x)+\delta_L, ..., \min(x) + (N_L-1)\delta_L, \tau, \tau+\delta_U, ..., \tau+(N_U-1)\delta_U, \max(x)$. In a similar fashion to the standard SV model, the centre of each interval, $x^1...x^N$, is defined as the mid-point as in equation 7.

Transition Probabilities

To condition the transition dynamics on the level of volatility it is necessary to compute the matrix of transition probabilities, $\widehat{q}^{i,j}$, such that it reflects the parameter values of the region to which x^j belongs. Based on the volatility dynamics given in equation 3,

$$\widehat{q}^{i,j} = \frac{\delta_i}{\sqrt{2\pi\sigma_{w,s}^2}} \exp\left(\frac{-\left(x^i - \alpha_s - \beta_s\, x^j\right)^2}{2\,\sigma_{w,s}^2}\right) \tag{22}$$

where $S = 1$ and $\delta_i = \delta_L\ \ if\ x^j < \tau$, otherwise $S = 2$ and $\delta_i = \delta_U$ when $x^i > \tau$.

The filtering algorithm proceeds as outlined earlier, recursing through equations 13 to 15 where the transition probabilities are now computed given equation 22. Finally, P_1^i is initialised from equation 17 using the ML estimates of the standard SV model.

3.3 Testing the significance of the size effect

Under the null hypothesis of $\theta_1 = \theta_2$ the SV model with size effects collapses to the standard SV model for any value of τ. Since τ is unidentified under this null, a standard likelihood ratio (LR) test will not follow a standard distribution. Therefore to obtain accurate inference regarding the adequacy of the size effect, it is proposed that the non-standard distribution of the LR statistic can be determined by the use of a bootstrap procedure.

This procedure can be summarised in the following steps:

1. Estimate the parameter vector, $\widehat{\theta}_{SV}$, of the standard SV model on actual data and store the log-likelihood, L_{SV}.
2. Estimate the parameter vector, $\widehat{\theta}_{SIZE}$, of the SV model with size effects on actual data and store the log-likelihood, L_{SIZE}.
3. Find the likelihood ratio statistic $LR = 2 \times (L_{SIZE} - L_{SV})$.
4. Set $i = 1$.
5. Simulate a return series of length T, from the standard SV process using the parameter vector $\widehat{\theta}_{SV}$.
6. Estimate the parameter vectors $\widehat{\theta}_{SV,i}$ and $\widehat{\theta}_{SIZE,i}$ storing the log likelihoods,$L_{SV,i}$ and $L_{SIZE,i}$.
7. Find the likelihood ratio statistic $LR_i = 2 \times (L_{SIZE,i} - L_{SV,i})$
8. Set $i = i + 1$ and repeat steps $5 - 7$ until $i = N_{sim}$.
9. The empirical $p - value$ is then found as $1/N_{sim} \sum_{i=1}^{N_{sim}} I_i$ where $I_i = 1$ if $LR_i > LR$ and 0 otherwise.

4 EMPIRICAL RESULTS

Two datasets are considered. Equity returns consisting of 2000 daily return observations from the S&P 500

Table 1. Estimation results for standard SV models and SV model with size effects (SIZE) for both the S&P500 and YEN/USD datasets. p-values for the LR statistic are generated from the bootstrap approach outlined above with Nsim = 500.

	$SIZE$			SV	
	$S\&P$	Yen		$S\&P$	Yen
α_1	$-1.1E-4$ $(3.3E-3)$	$-2.0E-3$ $(7.3E-3)$	α	-0.0059 $(3.8E-3)$	-0.011 $(5.1E-3)$
β_1	0.993 (5.3E−3)	0.983 (0.011)	β	0.9784 (7.9E−3)	0.966 (0.012)
$\sigma_{w,1}$	0.083 (0.012)	0.108 (0.024)	σ_w	0.1420 (0.022)	0.176 (0.031)
α_2	0.539 (0.614)	0.109 (0.178)			
β_2	0.426 (0.558)	0.832 (0.163)			
$\sigma_{w,2}$	0.457 (0.128)	0.485 (0.118)			
τ	0.238 (0.134)	0.133 (0.113)			
$Like$	−2660.9	−2599.4		−2667.6	−2609.8
LR	13.4	20.8			
p	0.002	< 0.002			

index spanning 5 September 1996 to 16 August 2004 are utilised. Currency returns in the form of 2000 daily YEN/USD observations spanning 29 November 1996 to 30 July 2004 are also considered. Both datasets have been standardised to zero mean and unit variance.

Parameter estimates for both the standard SV and SV with size effect, along with tests of significance are outlined in Table 1. As a benchmark, the results for the standard SV model are first addressed. These results reflect the commonly observed feature of relatively high persistence in conditional volatility. In comparison to these results, allowing for a size effect reveals a number of interesting features.

The most obvious result, once the size effect has been introduced is the difference between $\widehat{\theta}_1$and $\widehat{\theta}_2$ for both series. When volatility is relatively low ($< \tau$) it is more persistent than the standard SV case. Conversely, when volatility is quite high ($> \tau$) the persistence in volatility is much lower than the persistence found in either the low volatility region or the standard SV case. It is also evident that the variability of volatility is quite low (high) in the low (high) volatility regions.

In both cases the likelihood ratio tests indicate that the size effect is clearly statistically significant feature of the respective datasets. This implies that the conditional volatility of these series are not linear processes in that the dynamics of volatility is dependent upon the level of current volatility.

5 CONCLUSION

Much research attention has been paid to the dynamics of asset return volatility. Two significant features of volatility dynamics are of interest. These are the sign (level of volatility related to sign of past returns) and size (volatility dynamics related to current level of volatility) effects respectively. The asymmetric sign effect has been dealt with by numerous authors within both the GARCH and SV contexts. The size effect on the other hand has attracted much less attention with it not being considered in the context of an SV model.

The central contribution of this paper has been to propose a non-linear filtering based approach to the estimation of an SV process with size effects. A simple hypothesis testing procedure was also suggested to determine the significance of the size effect. While such a model has been considered here, the proposed DNF estimation procedure could be applied to a wider range of latent variable models where it is believed the the dynamics of the latent variable is related to its level. This has been achieved by partitioning the possible state-space into adjoining regions and utilising region specific transition probabilities within the prediction step within DNF algorithm.

Empirical application of the SV model with size effect shows that it is certainly an important feature of the two series considered here. Given the equity and currency returns considered, volatility dynamics appear to be dependent upon the current level of volatility. In both instances, the persistence of volatility falls and the volatility of volatility rises as the current level of volatility rises, suggesting that volatility dynamics are not linear.

6 REFERENCES

Andersen, T.G., 1996, "Return volatility and trading volume: an information flow interpretation of stochastic volatility", *Journal of Finance*, **51**, 169-204.

Bollerslev, T., 1987, "A Conditional Heteroskedastic time series model for speculative prices and rates of return", *Review of Economics and Statistics*, **69**, 542-547.

Bollerslev, T. and Chou, R.Y. and Kroner, K.F., 1992, "ARCH modelling in finance: a review of the theory and empirical evidence", *Journal of Econometrics*, **52**, 5-59.

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

Clark, P., 1973, "A subordinated stochastic process model with finite variance for speculative process", *Econometrica*, **41**, 135-155.

Clements A.E., Hurn A.S. and White S.I. (2004) "Discretised Non-Linear Filtering of Dynamic Latent Variable Models with Application to Stochastic Volatility" unpublished mimeo, Queensland University of Technology.

Ding, Z. and Granger, C.W.J. and Engle, R.F., 1993, "A Long Memory property of stock market Returns and a New Model", *Journal of Empirical Finance*, **1**, 83-105.

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

Engle, R. and Ng, V., 1993, "Measuring and Testing the Impact of News on Volatility", *Journal of Finance,* **48**, 1749-1778.

Friedman, B.M. and Laibson, D.I., 1989, "Economic implications of extroadinary movements in stock prices", *Brookings Papers on Economic Activity,* **2/89**, 137-189.

Ghysels, E. and Harvey, A.C. and Renault, E., 1996, "Stochastic volatility", in G.S. Maddala and C.R. Rao (eds), *Statistical Methods in Finance*, North Holland, Amsterdam.

Gouriéroux, C. and Monfort, A., 1992, "Qualitaive Threshold ARCH Models", *Journal of Econometrics,* **52**, 159-199.

Harvey, A. and Shephard, N., 1996, "Estimation of an asymmetric stochastic volatility model for asset returns", *Journal of business and Economic Statistics,* **14**, 429-434.

Hentschel, L., 1995, "All in the family: Nesting symmetric and assymetric garch models", *Journal of Financial Economics,* **39**, 71-104.

Kitagawa G. 1987. Non-gaussian state-space modeling of non-stationary time series. *Journal of the American Statistical Association* **82**: 1032-1062.

Longin, F., 1997, "The threshold effect in expected volatility: A model based on asymmetric information", *Review of Financial Studies*, **10**, 837-869.

MacDonald IL, Zucchini W. 1997. *Hidden Markov and Other Models for Discrete-valued Time Series*. Chapman and Hall: London.

Nelson, D., 1991, "Conditional heteroskedasticity in asset returns: a new approach", *Econometrica*, **59**, 347-370.

Pagan, A.R., 1996, "The econometrics of financial markets", *Journal of Empirical Finance,* **3**, 1, 15-102.

Shephard, N., 1996, "Statistical aspects of ARCH and stochastic volatility", in D. R. Cox, D. V. Hinkley and O. E. Barndorf-Nielsen (eds), *Time Series Models in Econometrics, Finance and Other Fields*, Chapman & Hall, London, 1-67.

So, M.K.P. and Li, W.K. and Lam K., 2002, "A threshold stochastic volatility model", *Journal of Forecasting,* **21**, 473-500.

Tauchen, G. and Pitts, M., 1983, "The price variability-volume relationship on speculative markets", *Econometrica*, **51**, 485-505.