
Stochastic volatility forecasting and risk management

Perry Sadorsky

*Schulich School of Business, York University, 4700 Keele Street,
Toronto, Ontario Canada M3J 1P3
Email: psadorsk@schulich.yorku.ca*

This paper compares the forecasting performance of the range-based stochastic volatility model with a number of other well-known forecasting models. Each forecasting model is applied to a financial data set that includes daily futures prices on, the S&P 500, ten year US government bond series, crude oil prices, and the foreign currency exchange rate between the Canadian and US dollar. Forecasts are evaluated out of sample using forecast summary statistics as well as value at risk measures like conditional coverage, independence and unconditional coverage. Overall the forecast summary statistics show that for each financial series, moving average, exponential smoothing and AR5 models to be better at forecasting the log range than the stochastic volatility model. Value at risk calculated from the stochastic volatility models does not reject independence in each of the four financial series studied but does reject conditional and unconditional coverage in all of the series studied. The empirical density model does not reject unconditional coverage in three out of the four financial series studied. All of the parametric models reject conditional coverage. These results show how difficult it is to design a good parametric value at risk model.

1. Introduction

Volatility modelling and forecasting in financial markets is an important and interesting topic to study. As financial markets around the world continue to move towards deregulation and globalization, the need for good forecasts of financial volatility as inputs into global risk management models continues to grow. Volatility forecasts are important inputs into option pricing models and risk management models. It is well known that many financial time series exhibit volatility clustering whereby volatility is likely to be high when it has recently been high and volatility is likely to be low when it has recently been low. Generalized autoregressive conditional heteroscedastic (GARCH) models are particularly useful for modelling time-varying

conditional volatility and GARCH models are extensively used by both researchers and practitioners. GARCH models the time varying variance as a deterministic function of lagged squared residuals and lagged conditional variance.

An alternative way to model time varying volatility is to use a stochastic volatility (SV) model (Taylor, 1982, 1986). SV models the variance as an unobserved component that follows a particular stochastic process. In this model, time-varying variance is not restricted to follow a deterministic process. In SV models it is usual to model volatility as a logarithmic first order autoregressive process. This is a discrete time approach to the diffusion process used in the option pricing literature (Hull and White, 1987). The SV model is theoretically attractive but empirically challenging because the unobserved

volatility process enters the model in a non-linear fashion which leads to the likelihood function depending upon high-dimensional integrals. A number of different approaches can be used to estimate stochastic volatility models.

Stochastic volatility models can be estimated using generalized method of moments (GMM) (Melino and Turnbull, 1990), quasi-maximum likelihood (QML) (Harvey *et al.*, 1994; Ruiz, 1994), efficient method of moments (EMM) (Gallant *et al.*, 1997; Andersen *et al.*, 1999), Markov chain Monte Carlo (MCMC) (Jacquier *et al.*, 1994; Kim *et al.*, 1998) and efficient importance sampling (EIS) (Richard and Zhang, 1996; Liesenfeld and Richard, 2003).

The various methods differ in ease and speed of estimation. MCMC and EIS dominate GMM and QML in small samples (less than 500 observations) but it is not clear whether the increased complexity of these models is justified in practice for larger samples. The QML method approximates a logarithmic chi-square process by a Gaussian process and uses the quasi-likelihood to approximate the likelihood function. QML is consistent and easy to estimate, but it is inefficient (although less so in large samples).

So *et al.* (1999) and Yu (2002) are two recent studies that compare the usefulness of the stochastic volatility model with GARCH models in applied forecasting situations. So *et al.* (1999) found that in modelling and forecasting foreign exchange rates, the stochastic volatility model estimated as a state space model does not, in general, outperform GARCH models. Yu (2002) uses the stochastic volatility model to forecast daily stock market volatility for New Zealand. Using forecast accuracy tests, he finds that the stochastic volatility model outperforms GARCH models. The mixed results from these two papers suggest that further research is needed on the relative merits of stochastic volatility models in applied forecasting situations.

Until recently, all of the stochastic volatility models have focused on using either squared returns or absolute returns as a proxy to the true but unobservable volatility. In a different approach to modelling stochastic volatility, Alizadeh *et al.* (2002) propose using the price range in the estimation of stochastic volatility models where the price range is defined as the difference between the highest and lowest log security prices over a fixed sampling interval. The information contained in the open, high, low, and close of a security price is widely used in Japanese candlestick charting techniques and other technical indicators (Nisson, 1991). Early academic work on the range based volatility estimator include Garman and Kless (1980), Parkinson

(1980), Beckers (1983) and Ball and Torous (1984). Alizadeh *et al.* (2002) suggest that the range is a better proxy to the true unobservable volatility because more information is contained in the range than in the squared daily returns. Security prices, for example, can fluctuate greatly over a trading day but it might happen that this day's closing price is similar to the previous day's closing price. In this case squared daily returns indicate little volatility. In comparison, the difference between the daily high and daily low would more accurately capture the daily volatility. Alizadeh *et al.* (2002) show that their range based volatility model is highly efficient and approximately Gaussian. They estimate and compare both one factor and two factor latent volatility models for foreign exchange futures prices and find that the two factor model has more desirable regression diagnostics.

Recent developments in risk management and financial engineering have highlighted the use of value at risk (VaR) as a popular approach to measuring market risk on a daily basis (JP Morgan, 1996; Jorion, 1997; Alexander, 2001; Brooks and Persaud, 2002, 2003; Christoffersen, 2003; Partnoy, 2003). Value at risk specifies the portfolio loss that could occur over a given time period with a given probability. For example, a 1% one day VaR is the loss that might occur one day in one hundred. VaR can be used to measure the risk in stocks, bonds, commodities, foreign currency exchange, options, futures, forwards, and swaps. In the parametric approach empirical models are used to estimate and forecast volatility $\hat{\sigma}_t$. VaR grew out of the work that the Basel Committee on Banking Supervision was doing on measuring market risk (1988 Basel Accord, 1996 Amendment). On 28 January 1997, the US Securities and Exchange Commission (SEC) passed a ruling requiring companies to disclose more information about their derivatives. VaR was one of the options that the SEC gave firms for disclosing the risks involved in trading derivatives. VaR measures are now included in the annual reports of companies doing business in the USA. In the parametric case, the VaR for a given level of significance α is, $\text{VaR}_\alpha = Z_\alpha \hat{\sigma}_t$, where Z_α is the appropriate constant from the standard normal tables. VaR can also be measured non-parametrically by using the empirical distribution to calculate the appropriate critical value corresponding to an α level of significance. Obviously, good forecasts of volatility are essential to getting good estimate of VaR.

This paper uses a range based stochastic volatility model to address a number of research questions. First, the forecasting performance of a large number of models, which include both regression and SV

models, are compared and contrasted for a set of data that is important from a global finance perspective. Second, the forecasting performance of models is compared and contrasted using a wider array of forecast statistics than commonly used in most other papers. Most papers evaluate the out-of-sample forecasting performance of models using standard forecast summary statistics like mean squared error, mean absolute deviation, mean percentage error, and Theil U statistics. In this paper forecast accuracy is also evaluated using regression based market timing equations. Third, forecasting performance is evaluated using value at risk measures.

This paper is organized as follows. In Section II the data are presented and analysed while in Section III the forecasting models are described. Forecast summary statistics are discussed in Section IV. Out of sample forecasting results are reported in Section V. Value at risk measures are presented and discussed in Section VI. Section VII concludes the paper.

II. Data

The data for this study consists of daily futures prices on the foreign currency exchange between the Canadian and US dollar, futures prices on the West Texas Intermediate crude oil contract, futures prices on the S&P 500 stock prices, and the futures prices on the US ten year Government note. All of the data are available from Prophet. The data set covers the period January 1984 to December 2003. The daily closing prices are transformed into continuously compounded rates of returns

$$r_t = 100 * \ln(S_t/S_{t-1})$$

where S_t is the closing price on day t and the sample size runs from 1 to T . These returns will be used to analyse value at risk. The range of the log-prices is defined as the difference between the daily log high price and the daily log low price.

$$D_t = \ln(S_t^H) - \ln(S_t^L)$$

The expected value of the squared range is

$$E[D_t^2] = 4\ln(2)\sigma^2$$

Thus, a proxy for the daily variance is

$$\sigma_{r,t}^2 = \frac{1}{4\ln(2)} D_t^2$$

For estimation purposes, Alizadeh *et al.* (2002) focus on the log range,

$$lD_t = \ln(D_t)$$

because the log range is approximately normally distributed in many applied situations. Consequently, in this paper, all models are estimated using the log range.

For each variable, the standard deviation is much smaller than its mean value and the mean value is very close to the median value (Table 1). The log range of oil prices and S&P 500 each display some evidence of skewness and/or kurtosis. The crude oil price series is skewed slightly to the left while the S&P series is skewed to the right. A normally distributed random variable has skewness of zero and kurtosis of three. The probability values for the Jarque and Bera (1980) test statistic indicate that the Canadian dollar series and the Treasury note series is each distributed normally.

Augmented Dickey and Fuller (1979) (ADF) and Phillips and Perron (1988) (PP) unit root tests for non-stationarity in the log range of the financial prices indicate no evidence of non-stationarity (Table 1). Each of the unit root test statistics is calculated with an intercept in the test regression. For each of these tests, the null hypothesis is a non-stationary time series and the alternative hypothesis is a stationary time series. The lag length for the ADF test regression is set using the Schwarz information criteria (SIC) and the bandwidth for the PP test regression is set using a Bartlett kernel.

Table 1. Summary statistics of log range

	<i>cd</i>	<i>cl</i>	<i>sp</i>	<i>ty</i>
Mean	-5.748	-3.896	-4.440	-5.272
Median	-5.737	-3.873	-4.450	-5.277
Maximum	-3.835	-1.261	-1.183	-3.310
Minimum	-7.503	-6.877	-6.194	-7.018
Std. Dev.	0.543	0.629	0.530	0.494
Skewness	-0.025	-0.201	0.306	0.081
Kurtosis	2.951	3.520	3.862	3.013
Jarque-Bera	1.024	90.016	234.012	5.449
Probability	0.599	0.000	0.000	0.066
Observations	5022	4992	5021	5006
Unit root tests				
ADF	0.000	0.000	0.000	0.000
PP	0.000	0.000	0.000	0.000

cd, *cl*, *sp*, and *ty* denote the futures prices on the foreign currency exchange between the Canadian and US dollar, futures prices on the West Texas Intermediate crude oil contract, futures prices on the S&P 500 stock prices, and the futures prices on the US ten year Government note. Probability values shown for unit root tests (Mackinnon, 1996). The lag length for the ADF test regression is set using the Schwarz information criteria (SIC) and the bandwidth (BW) for the PP test regressions is set using a Bartlett kernel.

Table 2. Autocorrelations of log range

Lags	<i>cd</i>	<i>cl</i>	<i>sp</i>	<i>ty</i>
1	0.496	0.629	0.485	0.307
2	0.454	0.602	0.491	0.289
3	0.423	0.587	0.458	0.262
4	0.412	0.571	0.459	0.275
5	0.412	0.568	0.441	0.309
6	0.391	0.548	0.429	0.238
7	0.389	0.534	0.413	0.241
8	0.384	0.528	0.424	0.231
9	0.376	0.527	0.413	0.24
10	0.384	0.526	0.409	0.227
11	0.383	0.511	0.391	0.228
12	0.364	0.514	0.379	0.193
s.e.e.	0.014	0.014	0.014	0.014

See Table 1 for variable definitions.

The first ten autocorrelations for the log range of each financial series are reported in Table 2. As a basis of comparison, recall that the autocorrelations for a randomly distributed variable should be less than two standard errors. The large and slowly decaying autocorrelations of the log range of each series show strong volatility persistence for each financial series.

III. Modelling and Forecasting Financial Price Volatility

This section provides a brief overview of the models used to forecast range based volatility. In this paper daily ex post volatility is measured by the squared range.

$$\sigma_{r,t}^2 = 0.361 D_t^2 \quad (1)$$

At time period t , a 1 day forecast is made. Models are estimated with five years of daily trading data for a total of 1250 observations (5×250).¹ The estimation period is then rolled forward by adding one new day and dropping the most distant day. In this way the sample size used in estimating the models stays at a fixed length and the forecasts do not overlap. Thus there are 3762 one day volatility forecasts for the Canadian dollar series, 3732 one day volatility forecasts for the oil prices series, 3761 one day volatility forecasts for the S&P 500 series, and 3746 one day volatility forecasts for the treasury note series.

The following models used are, random walk, historical mean, moving average, exponentially smoothing, linear regression model, autoregressive models, and a stochastic volatility model. Moving averages, exponential smoothing, linear regression and autoregressive models are fairly basic techniques in the applied forecasting literature (Hanke and Wichern, 2005) while the stochastic volatility model is at a much higher degree of complexity. All models are estimated using the log range as the response variable.

Random walk model

From a random walk (RW) model, the best forecast of next period's log range is this period's estimate of log range. As in other papers, the random walk model is used as the benchmark.

$$l\hat{D}_{t+1}(RW) = lD_t \quad (2)$$

Historical mean model

In the historical mean model, the best forecast of next period's log range is the average of the previous volatilities. This approach assumes a stationary log range series.

$$l\hat{D}_{t+1}(HM) = \frac{1}{1250} \sum_{j=0}^{1249} lD_{t-j} \quad (3)$$

Moving average model

Moving average (MA) methods are widely used in time series forecasting. In this paper a moving average of length m where $m = 20, 60, 180$ days is used to generate log range forecasts. Choosing these lengths is fairly standard because these values of m correspond to one month, three months and six months of trading days respectively. The expression for the m day moving average is shown below.

$$l\hat{D}_{t+1}(MA(m)) = \frac{1}{m} \sum_{j=0}^{m-1} lD_{t-j} \quad (4)$$

Exponential smoothing

Exponential smoothing (ES) models are also very widely used in applied forecasting. In ES models

¹The reported results for forecast summary statistics, value at risk, and their respective rankings are reasonably robust for small (one or two year) changes in the estimation sample size.

the current forecast of log range is calculated as the weighted average of the one period past value of log range and the one period past forecast of log range. This specification is appropriate provided the underlying log range series has no trend.

$$\hat{ID}_{t+1}(ES) = \alpha \hat{ID}_t(ES) + (1 - \alpha) ID_t \quad (5)$$

The smoothing parameter α lies between zero and unity. If α is zero then the ES model is the same as a random walk. If α is one then the ES model places all of the weight on the past forecast. In the estimation process the optimal value of α was chosen based on the root mean squared error criteria. The optimal α is the one that records the lowest MSE. The ES model and smoothing parameter are estimated for each forecast horizon using a 20 day, 60 day, 180, and 1250 day rolling window. Exponential smoothing is used to model volatility in JP Morgan's RiskMetrics (JP Morgan, 1996).

Least squares linear regression model

This model uses an ordinary least squares (OLS) regression model to model log range by using a one period lagged value of past log range as a driver.

$$\hat{ID}_{t,1}(LR) = \hat{\beta}_0 + \hat{\beta}_1 ID_{t-1} \quad (6)$$

Autoregressive model

This model uses an autoregressive (AR) process to model log range. Five lagged values (corresponding to one trading week) of past log range are used as drivers to make a one period ahead forecast.

$$\begin{aligned} \hat{ID}_{t+1}(AR5) = & \hat{\beta}_0 + \hat{\beta}_1 ID_t + \hat{\beta}_2 ID_{t-1} \\ & + \hat{\beta}_3 ID_{t-2} + \hat{\beta}_4 ID_{t-3} + \hat{\beta}_5 ID_{t-4} \end{aligned} \quad (7)$$

Discrete-time range-based stochastic volatility (SV) model

Alizadeh *et al.* (2002) present a formal derivation of the discrete time stochastic volatility model from the continuous time stochastic volatility model. In this paper, only the discrete time version is presented and the interested reader is referred to Alizadeh *et al.* (2002) for details.

The conditional distribution of log volatility is approximately

$$\ln \sigma_{(i+1)H} | \ln \sigma_{iH} \sim N[\ln \bar{\sigma} + \rho_H(\ln \sigma_{iH} - \ln \bar{\sigma}), \beta^2 H] \quad (8)$$

where $H = T/N$ and H is the sample length, T is the sample period and N are the number of intervals. The parameter β is a function that models the volatility of the unobserved volatility state variable v . It is convenient to work with a volatility proxy statistic that is a homogeneous in power γ .

$$f(s_{iH,(i+1)H}) = \sigma_{iH}^\gamma f(s_{iH,(i+1)H}^*) \quad (9)$$

Taking natural logarithms of both sides of Equation 9 yields,

$$\ln |f(s_{iH,(i+1)H})| = \gamma \ln \sigma_{iH} + \ln |f(s_{iH,(i+1)H}^*)| \quad (10)$$

where s^* denotes the continuous sample path of a standardized diffusion process.

The first term on the right-hand side of Equation 10 is proportional to log volatility and the second term is a measurement error. Following Harvey *et al.* (1994), Equations 8 and 10 can be written as a linear state space system and estimated using a Kalman filter.

$$\begin{aligned} \ln \sigma_{(i+1)H} = & \ln \bar{\sigma} + \rho_H(\ln \sigma_{iH} - \ln \bar{\sigma}) \\ & + \beta \sqrt{H} v_{(i+1)H} \end{aligned} \quad (11)$$

$$\begin{aligned} \ln |f(s_{iH,(i+1)H})| = & \gamma \ln \sigma_{iH} + E \ln |f(s_{iH,(i+1)H}^*)| \\ & + \varepsilon_{(i+1)H} \end{aligned} \quad (12)$$

Equation 11 is the state (transition) equation and Equation 12 is the signal equation. In Equation 12, E is the mathematical expectation operator. The state equation errors v are i.i.d. $N(0,1)$ and the signal equation errors have zero mean.

A two-factor model can be represented by the following state equation.

$$\ln \sigma_{(i+1)H} = \ln \bar{\sigma} + \ln \bar{\sigma}_{1,(i+1)H} + \ln \bar{\sigma}_{2,(i+1)H} \quad (13)$$

$$\ln \sigma_{1,(i+1)H} = \rho_{1,H} \ln \sigma_{1,iH} + \beta_1 \sqrt{H} v_{1,(i+1)H} \quad (14)$$

$$\ln \sigma_{2,(i+1)H} = \rho_{2,H} \ln \sigma_{2,iH} + \beta_2 \sqrt{H} v_{2,(i+1)H} \quad (15)$$

The error terms v_1 and v_2 are contemporaneously and serially independent $N(0,1)$ random variables.

IV. Forecast Summary Statistics

The forecast summary statistics include well known measures like mean squared error (MSE), mean absolute deviation (MAD), mean percentage error

(MPE), mean absolute percentage error (MAPE), and the Theil U statistic.

$$\begin{aligned} MSE &= \frac{\sum_{t=1}^n e_t^2}{n} \\ MAD &= \frac{\sum_{t=1}^n \text{abs}(e_t)}{n} \\ MPE &= \frac{\sum_{t=1}^n (e_t)/ID_t}{n} \\ MAPE &= \frac{\sum_{t=1}^n \text{abs}(e_t)/ID_t}{n} \\ U &= \frac{\sum_{t=1}^n (ID_{t+1} - \widehat{ID}_{t+1})}{\sum_{t=1}^n (ID_{t+1} - ID_t)} \end{aligned}$$

e_t is the forecast error between the actual and predicted values of the log range in time period t and n is the number of forecasts. These forecast summary statistics are useful for comparing models but they do not provide statistical tests of the difference between two models (Diebold, 1998, Chapter 12). For example, one model may have a lower MSE than another model but that doesn't mean that the difference between the two MSE values is statistically significant from zero. Consequently, the forecast summary statistics also include the Diebold and Mariano (1995) test for a mean differential loss function and several tests for market timing.

Diebold and Mariano (1995) develop a test of forecast accuracy between two sets of forecasts using the MSE. The null hypothesis of equal forecast accuracy is tested based on $E(d_t) = 0$ where E is the mathematical expectation operator and $d_t = e_{1t}^2 - e_{2t}^2$. The variables e_{1t} and e_{2t} are forecast errors from model 1 and model 2 respectively. The Diebold and Mariano (1995) test statistic is

$$DM = (\hat{V}(\bar{d}))^{-1/2} \bar{d} \sim N(0, 1)$$

where

$$\bar{d} = n^{-1} \sum_{t=1}^h d_t \quad \text{and} \quad \hat{V}(\bar{d}) \approx n^{-1} \left(\gamma_0 + 2 \sum_{k=1}^{h-1} \gamma_k \right)$$

where n is the number of forecasts computed from model 1 and model 2. The variable γ_k is the k th autocovariance of d_t . Under the null hypothesis, DM is asymptotically normally distributed. In this paper the DM test is calculated from a loss differential that compares the mean squared error of the corresponding model with the mean squared error of the random walk (Diebold, 1998, Chapter 12). A significant test statistic indicates that the selected forecasting model outperforms a random walk. Probability values are shown in the tables.

Harvey *et al.* (1997) propose a modified DM test (MDM) that improves upon the small sample properties of the DM test.

$$MDM = \left(\frac{n+1-2h+n^{-1}h(h-1)}{n} \right)^{1/2} DM$$

The MDM test for h -step ahead forecasts is distributed as a t distribution with $n-1$ degrees of freedom. Harvey *et al.* (1997) conduct Monte Carlo simulations to compare the performance of the MDM test with the DM test in a number of different applied situations. They find that the MDM test performs much better than the DM test across all forecast horizons and in situations where the forecast errors are autocorrelated or have non-normal distributions.

Another way of testing out-of-sample predictability is to use a forecast accuracy regression test (RT). This test compares the forecasted log range \widehat{ID} with the actual log range ID . This test provides a way to evaluate the ability of a forecasting model to predict both the correct direction and magnitude. The test is calculated by running the following ordinary least squares (OLS) regression.

$$ID_t = c_1 + c_2 \widehat{ID}_t + \xi_t$$

Newey and West (1987) heteroscedasticity and autocorrelation consistent probability values are shown in the tables. A good forecasting model should have no intercept (unbiased) and a slope of 1.

V. Out of Sample Forecasting Results

For the log range of the Canadian dollar series, the MSE ranks the AR5 first, MA20 second and ES third, while the MAD ranks the MA20 first, AR5 second and ES third (Table 3). The Theil U statistics ranks models in the same order as the MSE. The MSE and MAD each rank the two-factor stochastic volatility model ahead of the one-factor stochastic volatility model although both models are ranked towards the bottom. The MPE values indicate that every model over predicts and the ES20 has the least bias.

The regression based forecast accuracy tests (RT) indicate that, except for the MA180, all of the models have some bias (intercept significantly different from zero at the 5% level). All models have some forecasting power (the slope coefficient significantly different from zero at the 5% level). The RT test indicates that the MA180 is a good forecasting model because it has no bias and a slope coefficient statistically insignificant from one.

Table 3. Forecast summary statistics for log range of Canadian dollar

	RW	HM	MA20	MA60	MA180	ES	ES20	ES60	ES180	LR	AR5	SV1	SV2
MSE	0.304	0.252	0.191	0.205	0.223	0.192	0.199	0.195	0.194	0.214	0.191	0.219	0.216
MAD	0.440	0.402	0.347	0.359	0.374	0.348	0.354	0.352	0.350	0.367	0.347	0.372	0.363
MPE	-0.005	-0.024	-0.006	-0.007	-0.010	-0.006	-0.004	-0.005	-0.005	-0.017	-0.011	-0.020	-0.011
MAPE	-0.078	-0.073	-0.062	-0.064	-0.067	-0.062	-0.063	-0.063	-0.062	-0.066	-0.062	-0.067	-0.065
RANK MSE	13	12	2	7	11	3	6	5	4	8	1	10	9
RANK MAD	13	12	1	7	11	3	6	5	4	9	2	10	8
U		0.830	0.628	0.672	0.732	0.631	0.653	0.641	0.636	0.703	0.627	0.719	0.710
DM		0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MDM		0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RT													
c_1	-2.976	2.487	-0.565	-0.483	-0.241	-0.673	-1.010	-0.745	-0.712	0.587	0.295	2.184	-0.752
$p(c_1=0)$	0.000	0.000	0.000	0.000	0.136	0.000	0.000	0.000	0.000	0.000	0.027	0.000	0.009
c_2	0.477	1.415	0.901	0.914	0.955	0.881	0.824	0.870	0.875	1.092	1.046	1.367	0.863
$p(c_2=0)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$p(c_2=1)$	0.000	0.000	0.000	0.000	0.114	0.000	0.000	0.000	0.000	0.001	0.048	0.000	0.007

Probability values shown for DM and RT. The forecast summary statistics include mean squared error (MSE), mean absolute deviation (MAD), mean percentage error (MPE), mean absolute percentage error (MAPE), the Theil U statistic, Diebold and Mariano (1995) (DM); Harvey *et al.*, 1997 (MDM) test for a mean differential loss function (chosen model relative to a random walk), and a regression based forecast accuracy test (RT) test. The test is calculated by running an ordinary least squares (OLS) regression of actual volatility on forecasted volatility. The intercept coefficient is c_1 and the slope coefficient is c_2 . Newey and West (1987) heteroscedasticity and autocorrelation consistent probability values are shown in the tables. A good forecasting model should have no intercept (unbiased) and a slope of 1.

According to the DM and MDM tests, for each financial series, each of the forecasting models outperforms a random walk model (Tables 3, 4, 5, and 6).² This is encouraging because it means that the benchmark RW model is consistently beaten by each of the other models.

For the log range of the crude oil series, the MSE ranks the AR5 first, MA20 second and ES third, while the MAD ranks the MA20 first, ES second and AR5 third (Table 4). The MSE and MAD each rank the one factor stochastic volatility model ahead of the two factor stochastic volatility model although both models are ranked towards the bottom. The MPE values indicate that every model over predicts and the ES20 and ES60 have the least bias.

The regression based forecast accuracy tests (RT) indicate that, except for the AR5, all of the models have some bias (intercept significantly different from zero at the 5% level). All models have some forecasting power (the slope coefficient significantly different from zero at the 5% level).

For the log range of the S&P500 series, the MSE and the MAD rank the ES first, ES180 second and ES60 third (Table 5). The MSE and MAD each rank the two-factor stochastic volatility model ahead of the one-factor stochastic volatility model. The MPE values indicate that every model over predicts and the ES20 has the least bias.

The regression based forecast accuracy tests (RT) indicate that the LR, AR5 and SV1 each have no bias and each of these models has a slope coefficient statistically insignificantly different from 1. All models have some forecasting power (slope coefficient significantly different from zero at the 5% level). These results are important because even though the stochastic volatility model doesn't rank that high according to MSE or MAD, it is one of the three most preferred models according to RT.

For the log range of the US Treasury note series, the MSE ranks the AR5 first, ES second and ES180 third (Table 6). The MAD ranks the AR5 first, MA20 second and ES third. The MSE and MAD each rank the two factor stochastic volatility model ahead of the one-factor stochastic volatility model. The MPE values indicate that every model over predicts and the ES60 has the least bias.

The regression based forecast accuracy tests (RT) indicate that the LR and AR5 are each good forecasting models because they have no bias and each of these models has a slope coefficient statistically insignificant different from 1.

Overall these forecast summary statistics from Tables 3, 4, 5, and 6 shows that the AR5, moving average and exponential smoothing models to be better than the stochastic volatility models in every category (MSE, MAD, MPE, and MAPE). All models overpredict and the ES20 has the smallest bias for three of the four series studied. Based on the MSE, the AR5 model is the best prediction model in three out of the four securities studied.

The difference between the MSE for the best fitting and the best fitting SV model (either one-factor or two-factor) varies between 10% (treasury note) and 21% (oil). In comparison, the difference between MAD for the best fitting model and the best fitting SV model varies between 3.7% (Treasury note) and 9.3% (oil). Consequently, for three out of four series, the difference in MAD between the best fitting model and the SV model is not that great (less than 5%).

VI. Value at Risk

In this section value at risk measures are calculated and compared. In the parametric approach the empirical models discussed in previous sections are used to estimate and forecast volatility $\hat{\sigma}_t$. The VaR for a given level of significance α is, $VaR_\alpha = Z_\alpha \hat{\sigma}_t$, where Z_α is a constant from the standard normal tables.

Empirical density (ED) estimates of VaR were obtained by calculating the appropriate percentile (1%) of the actual sample of returns over the estimation period (1250 days).

VaR models can be backtested by first defining an indicator series (I) that takes on the value 1 (0) if the actual return in period $t+1$ is less (greater) than the forecasted VaR for the same time period.

$$I_{t+1} = 1, \quad \text{if } r_{t+1} < -VaR_{t+1} \\ I_{t+1} = 0, \quad \text{if } r_{t+1} > -VaR_{t+1}$$

The series I_{t+1} is known as the hit sequence (Christoffersen, 2003). Under the null hypothesis it should not be possible to predict whether the VaR will be violated in period $t+1$. Thus the hit sequence of violations will be distributed independently across time as a Bernoulli random variable with the following distribution function.

$$f(I_{t+1}, p) = (1 - p)^{1-I_{t+1}} p^{I_{t+1}}$$

p is the probability of the given level of significance at which the VaR is being calculated.

²In the analysis that follows, probability values from the DM and RT tests are being compared to 0.05.

Table 4. Forecast summary statistics for log range of crude oil

	RW	HM	MA20	MA60	MA180	ES	ES20	ES60	ES180	LR	AR5	SV1	SV2
MSE	0.289	0.269	0.181	0.194	0.239	0.181	0.183	0.182	0.181	0.207	0.178	0.215	0.248
MAD	0.430	0.404	0.334	0.347	0.383	0.334	0.338	0.336	0.335	0.360	0.334	0.365	0.382
MPE	-0.011	-0.038	-0.013	-0.014	-0.017	-0.013	-0.010	-0.010	-0.011	-0.025	-0.018	-0.031	-0.017
MAPE	-0.118	-0.114	-0.092	-0.096	-0.106	-0.092	-0.093	-0.092	-0.092	-0.100	-0.093	-0.102	-0.106
RANK	13	12	2	7	10	3	6	5	4	8	1	9	11
MSE													
RANK	13	12	1	7	11	2	6	5	4	8	3	9	10
MAD													
U		0.930	0.624	0.672	0.825	0.624	0.633	0.628	0.627	0.715	0.615	0.742	0.857
DM		0.023	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MDM		0.023	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RT													
$c1$	-1.985	-0.856	-0.403	-0.400	-0.943	-0.524	-0.673	-0.501	-0.477	-0.410	-0.160	1.127	-1.532
$p(c1=0)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.065	0.000	0.000
$c2$	0.476	0.761	0.894	0.894	0.751	0.862	0.824	0.870	0.876	0.883	0.953	1.280	0.594
$p(c2=0)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$p(c2=1)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.040	0.000	0.000

See Table 3.

Table 5. Forecast summary statistics for log range of S&P 500

	RW	HM	MA20	MA60	MA180	ES	ES20	ES60	ES180	LR	AR5	SV1	SV2
MSE	0.287	0.268	0.174	0.189	0.205	0.170	0.178	0.173	0.171	0.214	0.175	0.226	0.195
MAD	0.432	0.413	0.330	0.345	0.359	0.325	0.334	0.330	0.327	0.370	0.332	0.380	0.339
MPE	-0.008	-0.019	-0.009	-0.009	-0.010	-0.009	-0.007	-0.009	-0.008	-0.014	-0.012	-0.017	-0.010
MAPE	-0.099	-0.097	-0.076	-0.080	-0.084	-0.075	-0.077	-0.076	-0.076	-0.086	-0.077	-0.088	-0.078
RANK MSE	13	12	4	7	9	1	6	3	2	10	5	11	8
RANK MAD	13	12	4	8	9	1	6	3	2	10	5	11	7
U		0.932	0.607	0.658	0.713	0.593	0.619	0.603	0.595	0.744	0.609	0.787	0.679
DM		0.019	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MDM		0.019	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RT													
$c1$	-2.233	-1.222	-0.419	-0.395	-0.382	-0.409	-0.685	-0.480	-0.372	-0.073	0.117	0.093	-0.805
$P(c1=0)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.549	0.207	0.506	0.000
$c2$	0.499	0.722	0.906	0.912	0.915	0.908	0.847	0.892	0.917	0.981	1.024	1.016	0.818
$P(c2=0)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$P(c2=1)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.483	0.247	0.615	0.000

See Table 3.

Table 6. Forecast summary statistics for log range of US 10 year Treasury note

	RW	HM	MA20	MA60	MA180	ES	ES20	ES60	ES180	LR	AR5	SV1	SV2
MSE	0.345	0.233	0.201	0.210	0.219	0.200	0.209	0.203	0.201	0.218	0.198	0.220	0.218
MAD	0.465	0.386	0.356	0.365	0.372	0.356	0.363	0.359	0.357	0.373	0.355	0.375	0.368
MPE	-0.006	-0.011	-0.007	-0.008	-0.010	-0.007	-0.006	-0.005	-0.006	-0.010	-0.009	-0.010	-0.010
MAPE	-0.089	-0.074	-0.068	-0.070	-0.072	-0.068	-0.070	-0.069	-0.069	-0.072	-0.068	-0.072	-0.071
RANK MSE	13	12	4	7	10	2	6	5	3	8	1	11	9
RANK MAD	13	12	2	7	9	3	6	5	4	10	1	11	8
U		0.675	0.583	0.608	0.634	0.580	0.604	0.587	0.581	0.631	0.574	0.637	0.632
DM		0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MDM		0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RT													
$c1$	-3.920	-2.226	-1.193	-1.095	-0.836	-1.145	-1.772	-1.312	-1.048	-0.405	-0.074	1.206	-1.873
$P(c1=0)$	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000	0.000	0.191	0.708	0.005	0.000
$c2$	0.261	0.579	0.775	0.793	0.841	0.784	0.667	0.754	0.803	0.922	0.985	1.224	0.645
$P(c2=0)$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$P(c2=1)$	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000	0.000	0.178	0.680	0.006	0.000

See Table 3.

The unconditional coverage hypothesis can be used to test if the fraction of violations (π) from a particular risk model is significantly different from the theoretical fraction p . The likelihood of an independent and identically distributed Bernoulli hit sequence is

$$L(\pi) = \prod_{t=1}^T (1 - \pi)^{1-I_{t+1}} \pi^{I_{t+1}} = (1 - \pi)^{T_0} \pi^{T_1}$$

The variables T_0 and T_1 are the number of zeros and ones in the hit sequence. π is estimated from (T_1/T) . Under the null hypothesis, $\pi = p$ and under the alternative hypothesis, $\pi = (T_1/T)$. The unconditional coverage hypothesis can be tested using a likelihood ratio test.

$$LR_{uc} = -2 \ln[L(p)/L(\hat{\pi})] \sim \chi_1^2$$

Theoretically, VaR violations should be independent of each other. In actual practice VaR violations often show clustering. This is especially the case if the VaR was estimated from historical averages. A test for independence can be based on the assumption that a hit sequence that is dependent across time can be described by a Markov transition probability matrix.

$$\Pi_1 = \begin{bmatrix} 1 - \pi_{01} & \pi_{01} \\ 1 - \pi_{11} & \pi_{11} \end{bmatrix}$$

π_{11} is the probability of tomorrow being a violation conditional on today being a violation. π_{01} is the probability of tomorrow being a violation conditional on today being a non-violation. Empirically, π_{11} and π_{01} can be estimated from

$$\hat{\pi}_{01} = \frac{T_{01}}{T_{00} + T_{01}}$$

$$\hat{\pi}_{11} = \frac{T_{11}}{T_{10} + T_{11}}$$

Probabilities sum to one and therefore

$$\hat{\pi}_{00} = 1 - \hat{\pi}_{01}$$

$$\hat{\pi}_{10} = 1 - \hat{\pi}_{11}$$

The independence hypothesis ($\pi_{01} = \pi_{11}$) can be tested using a likelihood ratio test.

$$LR_{ind} = -2 \ln[L(\hat{\pi})/L(\hat{\Pi}_1)] \sim \chi_1^2$$

A joint test of conditional coverage (LR_{cc}) can be calculated by summing the test for unconditional coverage with the test for independence (Christoffersen, 2003).

For the Canadian dollar volatility series, unconditional coverage is rejected at the 5% level for each model except the ED model (Table 7, first panel). The hypothesis of independence in the hit sequence is

rejected at the 5% level for the RW, MA20, ES, ES20, ES60 and ED models. Conditional coverage is rejected at the 5% level of significance by all of the models.

For the crude oil volatility series, unconditional coverage is rejected at the 5% level for each model except the ED model (Table 7, second panel). The hypothesis of independence in the hit sequence is rejected at the 5% level for the HM, MA20, MA60, ES, and ED models. Conditional coverage is rejected at the 5% level of significance by all of the models.

For the S&P 500 volatility series, unconditional coverage is rejected at the 5% level for each model (Table 7, third panel). The hypothesis of independence in the hit sequence is rejected at the 5% level for the RW, HM, MA60, MA180, and ED models. Conditional coverage is rejected at the 5% level of significance by all of the models.

For the US Treasury note volatility series, unconditional coverage is rejected at the 5% level for each model except the ED (Table 7, fourth panel). The hypothesis of independence in the hit sequence is not rejected at the 5% level by any model. Conditional coverage is rejected at the 5% level of significance by all of the models.

VaR calculated from the stochastic volatility models does not reject independence in each of the four financial series studied but does reject conditional and unconditional coverage in all of the series studied. In general, most of the models studied do not reject independence but do reject conditional and unconditional coverage and all of the parametric models reject conditional coverage. Like Brooks and Persaud (2002) these results show how difficult it is to design a good parametric value at risk model.

VaR calculated from the stochastic volatility models does not reject independence in each of the four financial series studied and VaR calculated from the empirical density function does not reject unconditional coverage in three out of the four financial series studied. This suggests that it might be useful to combine the VaR from a stochastic volatility model with a VaR from the empirical density function. In this way a new VaR measure is constructed that satisfies both unconditional coverage and independence (and in doing so will most likely satisfy conditional coverage).

VII. Conclusions

This paper compares the forecasting performance of the range-based stochastic volatility model with a number of other well-known forecasting models. The information contained in the high and low of

Table 7. Value at Risk summary statistics

	RW	HM	MA20	MA60	MA180	ES	ES20	ES60	ES180	LR	AR5	SV1	SV2	ED
<i>cd</i>														
LR_{uc}	714.405	814.184	511.388	534.668	634.913	519.113	503.698	519.113	511.388	748.681	622.616	770.339	748.681	2.666
LR_{ind}	4.700	0.785	5.281	3.368	2.522	3.942	5.638	6.031	3.316	0.928	1.564	0.000	0.928	5.003
LR_{cc}	719.105	814.969	516.669	538.036	637.435	523.056	509.336	525.144	514.704	749.609	624.180	770.339	749.609	7.670
$p-LR_{uc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.102
$p-LR_{ind}$	0.030	0.376	0.022	0.066	0.112	0.047	0.018	0.014	0.069	0.335	0.211	0.990	0.335	0.025
$p-LR_{cc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.006
<i>cl</i>														
LR_{uc}	469.016	680.554	454.037	450.315	522.586	413.623	424.530	399.218	395.641	557.931	457.768	565.878	557.931	0.836
LR_{ind}	2.637	9.385	4.600	5.913	3.799	4.213	1.301	3.764	1.337	0.138	0.294	1.129	0.138	10.296
LR_{cc}	471.653	689.939	458.637	456.228	526.384	417.835	425.830	402.982	396.978	558.069	458.062	567.007	558.069	11.131
$p-LR_{uc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.361
$p-LR_{ind}$	0.104	0.002	0.032	0.015	0.051	0.040	0.254	0.052	0.247	0.710	0.588	0.288	0.710	0.001
$p-LR_{cc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001
<i>sp</i>														
LR_{uc}	354.353	399.975	205.143	268.559	256.052	213.853	216.784	210.936	213.853	303.929	252.956	337.273	303.929	5.652
LR_{ind}	4.193	7.376	0.017	6.666	5.925	0.053	0.021	0.046	0.032	0.075	0.235	0.330	0.075	7.334
LR_{cc}	358.546	407.350	205.159	275.224	261.977	213.906	216.805	210.982	213.886	304.004	253.191	337.603	304.004	12.986
$p-LR_{uc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.017
$p-LR_{ind}$	0.041	0.007	0.898	0.010	0.015	0.818	0.885	0.831	0.858	0.784	0.628	0.566	0.784	0.007
$p-LR_{cc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
<i>ty</i>														
LR_{uc}	373.088	254.046	208.997	244.806	266.537	217.773	223.690	206.098	203.213	244.806	217.773	238.709	244.806	3.844
LR_{ind}	3.331	3.350	0.056	0.044	0.467	2.607	1.491	0.327	1.000	0.030	0.018	0.078	0.030	0.148
LR_{cc}	376.418	257.395	209.053	244.850	267.004	220.380	225.181	206.425	204.214	244.836	217.791	238.787	244.836	3.992
$p-LR_{uc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.050
$p-LR_{ind}$	0.068	0.067	0.812	0.834	0.494	0.106	0.222	0.567	0.317	0.863	0.893	0.780	0.863	0.701
$p-LR_{cc}$	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.046

Probability values are shown for likelihood ratio (LR) tests of unconditional coverage (uc), independence (ind) and conditional coverage (cc).

a security's price is very often used by momentum traders to help them evaluate trading possibilities. The price range is a natural estimate to the true but unobserved volatility. Until recently, however, most academic research on volatility forecasting has used squared daily returns as a proxy for daily volatility. Each forecasting model is applied to a daily financial futures price data set that includes the S&P 500, ten year US government bond series, crude oil prices, and the foreign currency exchange rate between the Canadian and US dollar.

Overall the out of sample forecast summary statistics show that for each financial series, moving average, exponential smoothing and AR5 models to be better at forecasting the log range than the stochastic volatility model. This is useful information because these models are considerably easier to estimate than stochastic volatility models. The difference between MAD for the best fitting model and the best fitting SV model varies between 3.7% (Treasury note) and 9.3% (oil). Consequently, for three out of four series, the difference in MAD between the best fitting model and the SV model is not that great (less than 5%).

Value at risk measures are evaluated using unconditional coverage, independence and conditional coverage. Value at risk calculated from the stochastic volatility model does not reject independence in each of the four financial series studied but does reject conditional and unconditional coverage in all of the series studied. All of the parametric models reject conditional coverage. These results show how difficult it is to design a good parametric value at risk model.

Acknowledgements

Parts of this paper were presented at the 11th Annual Global Finance Conference in Las Vegas. I thank participants for their useful comments.

References

Alexander, C. (2001) *Market Models: A Guide to Financial Data Analysis*, John Wiley and Sons, New York.

Alizadeh, S. Brandt, M. W. and Diebold, F. X. (2002) Range-based estimation of stochastic volatility models, *Journal of Finance*, **LV11**, 1047–91.

Andersen, T., Chung, H. and Sorensen, B. E. (1999) Efficient method of moments estimation of a stochastic volatility model: a Monte Carlo study, *Journal of Econometrics*, **91**, 61–87.

Ball, C. A. and Torous, W. N. (1984) The maximum likelihood estimation of security price volatility: theory, evidence, and application to option pricing, *Journal of Business*, **57**, 97–112.

Beckers, S. (1983) Variance of security price returns based on high, low and closing prices, *Journal of Business*, **56**, 97–112.

Brooks, C. and Persaud, G. (2002) Model choice and Value-at-Risk performance, *Financial Analysts Journal*, **58**(5), 87–97.

Brooks, C. and Persaud, G. (2003) Volatility forecasting for risk management, *Journal of Forecasting*, **22**, 1–22.

Christoffersen, P. F. (2003) *Elements of Financial Risk Management*, Elsevier Science, USA.

Dickey, D. and Fuller, W. (1979) Distribution of the estimators for autoregressive time series with a unit root, *Journal of the American Statistical Association*, **74**, 427–31.

Diebold, F. X. (1998) *Elements of Forecasting*, South Western, Cincinnati, OH.

Diebold, F. X. and Mariano, R. (1995) Comparing predictive accuracy, *Journal of Business and Economic Statistics*, **13**, 253–65.

Gallant, A. R., Hsieh, D. A. and Tauchen, G. E. (1997) Estimation of stochastic volatility models with diagnostics, *Journal of Econometrics*, **81**, 159–92.

Garman, M. B. and Klass, M. J. (1980) On the estimation of price volatility from historical data, *Journal of Business*, **53**, 67–78.

Hanke, J. E. and Wichern, D. W. (2005) *Business Forecasting*, 8th edn, Prentice Hall, New York.

Harvey, A. C., Ruiz, E. and Shephard, N. (1994) Multivariate stochastic variance models, *Review of Economics Studies*, **61**, 247–64.

Harvey, D., Leybourne, S. and Newbold, P. (1997) Testing the equality of prediction mean squared errors, *International Journal of Forecasting*, **13**, 281–91.

Hull, J. and White, A. (1987) The pricing of options on assets with stochastic volatilities, *Journal of Finance*, **42**, 281–300.

Jacquier, E., Polson, N. G. and Rossi, P. E. (1994) Bayesian analysis of stochastic volatility models (with discussion), *Journal of Business and Economic Statistics*, **12**, 371–89.

Jarque, C. M. and Bera, A. K. (1980) Efficient tests for normality, homoskedasticity and serial independence of regression residuals, *Economics Letters*, **6**, 225–59.

Jorion, P. (1997) *Value at Risk: The New Benchmark for Controlling Market Risk*, McGraw-Hill, New York.

JP Morgan (1996) *RiskMetrics—Technical Document*, 4th edn, JP Morgan, New York.

Kim, S., Shephard, N. and Chib, S. (1998) Stochastic volatility: likelihood inference and comparison with ARCH models, *Review of Economics Studies*, **65**, 361–93.

Liesenfeld, R. and Richard, J. F. (2003) Univariate and multivariate stochastic volatility models: estimation and diagnostics, *Journal of Empirical Finance*, **10**, 505–31.

Mackinnon, J. G. (1996) Numerical distribution functions for unit root and cointegration tests, *Journal of Applied Econometrics*, **11**, 601–18.

Melino, A. and Turnbull, S. (1990) Pricing foreign currency options with stochastic volatility, *Journal of Econometrics*, **45**, 239–65.

Newey, W. and West, K. (1987) A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica*, **55**, 703–8.

- Nison, S. (1991) *Japanese Candlestick Charting Techniques*, New York Institute of Finance, New York.
- Parkinson, M. (1980) The extreme value method for estimating the variance of the rate of return, *Journal of Business*, **53**, 61–5.
- Partnoy, F. (2003) *Infectious Greed: How Deceit and Risk Corrupted the Financial Markets*, Owl Books, Henry Holt and Company, New York.
- Phillips, P. C. B. and Perron, P. (1988) Testing for a unit root in time series regression, *Biometrika*, **75**, 335–46.
- Richard, J. F. and Zhang, W. (1996) Econometric modelling of UK house prices using accelerated importance sampling, *Oxford Bulletin of Economics and Statistics*, **58**, 601–13.
- Ruiz, E. (1994) Quasi-maximum likelihood estimation of stochastic volatility models, *Journal of Econometrics*, **63**, 289–306.
- So, M. K. P., Lam, K. and Li, W. K. (1999) Forecasting exchange rate volatility using autoregressive random variance model, *Applied Financial Economics*, **9**, 583–91.
- Taylor, S. J. (1982) Financial returns modeled by the product of two stochastic processes – a study of daily sugar prices, in *Time Series Analysis: Theory and Practice 1*, (Ed.) O. D. Anderson, North Holland, Amsterdam, pp. 203–26.
- Taylor, S. J. (1986) *Modelling Financial Time Series*, Wiley, Chichester.
- Yu, J. (2002) Forecasting volatility in the New Zealand stock market, *Applied Financial Economics*, **12**, 193–202.