Stochastic Volatility Models Predictive Relevance for Equity Markets

Per Bjarte Solibakke
[1] Norwegian University of Science and Technology, Faculty of Economics and Management
[2] Larsgårdsvn. 2, 6025 Ålesund, Norway
per.b.solibakke@ntnu.no

Abstract. This paper builds and implements multifactor stochastic volatility models where the main objective is step ahead volatility prediction and to describe its relevance for the equity markets. The paper outlines stylised facts from the volatility literature showing density tails, persistence, mean reversion, asymmetry and long memory, all contributing to systematic data dependencies. As a by-product of the multifactor stochastic volatility model estimation, a long-simulated realization of the state vectors is available. The realization establishes a functional form of the conditional distribution, which is evaluated on observed data convenient for step ahead predictions. The paper uses European equity for relevance arguments and illustrational prediction purposes. Multifactor SV models empower volatility visibility and predictability enriching the amount of information available for equity market participants.

Keywords: Stochastic Volatility, Markov Chain Monte Carlo (MCMC) Simulations, Projection-Reprojection

1 Introduction

This paper builds and assesses multifactor scientific stochastic volatility (SV) models for the prediction of equity market volatility. Volatility is a measure of dispersion around the mean return of an asset. When the price returns are tightly bunched together (or spread apart), the volatility is small (large). The use of all volatility models entails prediction characteristics for future returns. A volatility model has been used internationally to predict the absolute magnitude of returns, quantiles and entire densities. A special feature of asset volatility is that it is not directly observable. The unobservability of volatility makes it difficult to evaluate the forecasting performance of volatility models. However, knowledge of the empirical properties of future prices is important when constructing risk management strategies, i.e.d portfolio selection, derivatives and hedging, market making and market timing. For all these activities, the predictability of volatility is essential for success. Modern portfolio theory (MPT) suggests that volatility creates risk. Portfolio studies have shown that when volatility increases, risk increases, and portfolio returns decreases. An equity risk manager therefore would want to know the likelihood of future asset and portfolio movements. If a portfolio manager adds more assets to his portfolio, the additional assets diversify the portfolio if they do not covary (correlation less than 1) with other assets in the portfolio. Hence, generally, portfolios imply risk reduction through diversification suggesting asset allocation importance. Mean-variance analysis and the Capital Asset Pricing Model are natural extensions of the portfolio analysis. An equity derivative trader wants to know the volatility that can be expected as contracts mature for both pricing and general risk management activities. The most important use of derivatives is a risk-reduction technique known as hedging, which requires a sound understanding of how to value derivatives and an understanding of which risks should and should not be hedged. Generally, for hedging, an equity risk manager will want to know the contract volatility approaching maturity. The only parameter that requires estimation in the Black-Scholes Model is the volatility. This volatility estimate also may be of use in estimating parameters (u and d) in a binomial model. Ceteris paribus, higher (lower) volatility increases (decreases) derivative prices. Therefore, market participants will sell (buy) both call and put option contract positions that are not part of speculative or hedge positions, if predicted volatility is declining (increasing). In contrast, a portfolio manager may want to buy (sell) a stock or a portfolio before its volatility falls (rises). Finally, a market maker can change his bid-ask spread believing future volatility changes. Normally, the equity markets show that the bid-ask spread increases (decreases) when volatility rises (falls).

Stochastic volatility models have an intuitive and simple structure and can explain the major stylized facts of asset, currency and commodity price changes. The motivation for stochastic volatility is the observed non-constant and frequently changing volatility. Time-varying volatility is endemic in financial markets and market participants who understand the dynamic behaviour of volatility are more likely to have realistic expectations about future prices and the risks to which they are exposed. The SV implementation is an attempt to specify how the volatility changes over time. Bearing in mind that volatility is a non-traded instrument, which suggests imperfect estimates, the volatility can be interpreted as a latent variable that can be modelled and predicted through its direct influence on the magnitude of returns. Risks may change through time in complicated ways, and it is natural to build multifactor stochastic models for the temporal evolution in volatility. The implementation adapts the MCMC estimator proposed by Chernozhukov and Hong [9], claimed to be substantially superior to conventional derivative based hill climbing optimizers for this stochastic class of problems. Moreover, under correct specification of the structural models the normalized value of the objective function is asymptotically χ^2 distributed (and the degrees of freedom is specified). The paper focuses on the Bayesian Markov Chain Monte Carlo (MCMC) modelling strategy used by Gallant and McCulloch [18] and Gallant and Tauchen[1] [14], [19] implementing multivariate statistical models derived from scientific considerations. The method is a systematic approach to generate moment conditions for the generalized method of moments (GMM) estimator [24] of the parameters of a structural model. Moreover, the implemented Chernozhukov and Hong [9] estimator keeps model

---

[1] The methodology is designed for estimation and inference for models where (1) the likelihood is not available, (2) some variables are latent (unobservable), (3) the variables can be simulated and (4) there exist a well-specified and adequate statistical model for the simulations. The methodologies (General Scientific Models (GSM) and Efficient Method of Moments (EMM)) are general-purpose implementation of the Chernozhukov and Hong [9] estimator.

parameters in the region where predicted shares are positive for every observed price/expenditure vector. Moreover, the methodology supports restrictions, inequality restrictions, and informative prior information (on model parameters and functionals of the model). This article is organized as follows. Section 2 describes the SV methodology. Section 3 presents stylized facts and section 4 concretizes these facts from stochastic volatility models showing two examples, one index and one asset. Section 5 summarizes and concludes.

2 Theory and Methodology

2.1 Stochastic Volatility Models

The SV approach specifies the predictive distribution of price returns indirectly, via the structure of the model, rather than directly. The SV model has its own stochastic process without worries about the implied one-step-ahead distribution of returns recorded over an arbitrary time interval convenient for the econometrician. The starting point is the application of Andersen et al. [2] considering the familiar stochastic volatility diffusion for an observed stock price S_t given by

$$\frac{dS_t}{S_t} = \left(\mu + c\left(V_{1,t} + V_{2,t}\right)\right)dt + \sqrt{V_{1,t}}\,dW_{1,t} + \sqrt{V_{2,t}}\,dW_{2,t} \tag{1}$$

where the unobserved volatility processes $V_{i,t}$, i = 1,2, is either log linear or square root (affine). The $W_{1,t}$ and $W_{2,t}$ are standard Brownian motions that are possibly correlated with corr($dW_{1,t}$, $dW_{2,t}$) = ρ. Andersen et al. [2] estimate both versions of the stochastic volatility model with daily S&P500 stock index data, 1953-December 31, 1996. Both SV model versions are sharply rejected. However, adding a jump component to a basic SV model greatly improves the fit, reflecting two familiar characteristics: thick non-Gaussian tails and persistent time-varying volatility. A SV model with two stochastic volatility factors show encouraging results in Chernov et al. [8]. The authors consider two broad classes of setups for the volatility index functions and factor dynamics: an affine setup and a logarithmic setup. The models are estimated using daily data on the Dow Index, January 2, 1953-July 16, 1999. They find that models with two volatility factors do much better than do models with only a single volatility factor. They also find that the logarithmic two-volatility factor models outperform affine jump diffusion models and provide acceptable fit to the data. One of the volatility factors is extremely persistent and the other strongly mean reverting.

This paper's SV model applies the logarithmic model with two stochastic volatility factors [8]. The model is extended to facilitate correlation between the mean and the stochastic volatility factors. The correlation applies the Cholesky decomposition for consistence. The main argument for the correlation modelling is to introduce asymmetry effects (correlation between return innovations and volatility innovations). The formulation of a general SV model for price change processes (y_t) therefore becomes

$$
\begin{aligned}
y_t &= a_0 + a_1\left(y_{t-1} - a_0\right) + \exp(V_{1t} + V_{2t}) \cdot u_{1t} \\
V_{1t} &= b_0 + b_1\left(V_{1,t-1} - b_0\right) + u_{2t} \\
V_{2t} &= c_0 + c_1\left(V_{2,t-1} - c_0\right) + u_{3t} \\
u_{1t} &= dW_{1t} \\
u_{2t} &= s_1\left(r_1 \cdot dW_{1t} + \sqrt{1-r_1^2} \cdot dW_{2t}\right) \\
u_{3t} &= s_2\left(\begin{array}{l} r_2 \cdot dW_{1t} + \left((r_3 - (r_2 \cdot r_1)) / \sqrt{1-r_1^2}\right) \cdot dW_{2t} + \\ \sqrt{1-r_2^2 - \left((r_3 - (r_2 \cdot r_1)) / \sqrt{1-r_1^2}\right)^2} \cdot dW_{3t} \end{array}\right)
\end{aligned}
\tag{2}
$$

where $W_{i.t}, i$ =1,2 and 3 are standard Brownian motions (random variables). The parameter vector is θ. The r's are correlation coefficients from a Cholesky decomposition[2]; enforcing an internally consistent variance/covariance matrix. Early references are Rosenberg [32], Clark [10], Taylor [36] and Tauchen and Pitts [35]. References that are more recent are Gallant et al. [15], [18], [20], Andersen [1], Durham [12], Shephard [34], Taylor [37], and Chernov et al. [8]. The model above has three stochastic factors. Even jumps with the use of Poisson distributions for jump intensity are applicable (complicates estimations considerably). The paper applies a computational methodology proposed by Gallant and McCulloch [17] and Gallant and Tauchen [19], [20] for statistical analysis of a stochastic volatility model derived from a scientific process[3]. Intuitively, the approach may be explained as follows. First, a reduced-form auxiliary model is estimated to have a tractable likelihood function (generous parameterization). The estimated set of score moment functions encodes important information regarding the probabilistic structure of the raw data sample. Second, a long sample is simulated from the continuous time SV model. Using the Metropolis-Hastings algorithm and parallel computing, parameters are varied in order to produce the best possible fit to the quasi-score moment functions evaluated on the simulated data. An extensive set of model diagnostics and an explicit metric for measuring the extent of SV model failure are useful side-products. The scientific stochastic volatility model cannot generate likelihoods, but it can be easily simulated.

2.2 The unobserved state vector using the nonlinear Kalman filter

From the prior SV model estimation, a by-product is a long simulated realization of the state vector $\{\hat{V}_{i,t}\}_{t=1}^{N}, i=1,2$ and the corresponding $\{\hat{y}_t\}_{t=1}^{N}$ for $\theta = \hat{\theta}$. Hence, by calibrating the functional form of the conditional distribution of functions given $\{\hat{y}_\tau\}_{\tau=1}^{t}$; evaluating

---

[2] For the Cholesky decomposition methodology see [4]

[3] See www.econ.duke.edu/webfiles/arg for software and applications of the MCMC Bayesian methodology. All models are coded in C/C++ and executable in both serial and parallel versions (OpenMPI).

the result on observed data $\{\tilde{y}_t\}_{t=1}^{n}$; generating predictions for $V_{i,t}, i=1,2$ through Kalman filtering y_t, very general functions of $\{y_\tau\}_{\tau=1}^{t}$ can be used and a huge dataset is available. An SNP model is estimated on the $\hat{y}_t$. The model represents a one-step ahead conditional variance $\hat{\sigma}_t^2$ of $\hat{y}_{t+1}$ given $\{\tilde{y}_\tau\}_{\tau=1}^{t}$. Regressions are run of $\hat{V}_{i,t}$ on $\hat{\sigma}_t^2$, $\hat{y}_t$ and $|\hat{y}_t|$ and lags (generously long) of these series. These functions are evaluated on the observed data series $\{\tilde{y}_\tau\}_{\tau=1}^{t}$, which give values $\tilde{V}_{i,t}, i=1,2$ for the volatility factors at the original data points.

3 Stylized facts of volatility

Modelling and forecasting market volatility have been the subject of vast empirical and theoretical investigation over the past two decades by academics and practitioners. Volatility, as measured by the standard deviation or variance of returns, is often used as a crude measure of total risk. The volatility is not directly observable making it difficult to evaluate the forecasting performance. A good volatility model must be able to capture and reflect the stylized facts. Moreover, a good volatility model should predict volatility for success. The task of forecasting volatility conditional on previously observed data is akin to filtering in Markov-Chained Monte-Carlo (MCMC) analyses[4]. Eliciting dynamics from observables are the one-step-ahead conditional volatility $Var(y_0 \mid x_{-1})$, where $x_{-1} = (y_{-L},...,y_{-1})$. The volatility can be obtained from standard recursions for the moments of the normal [26]. Filtered volatility is one-step-ahead conditional standard deviation evaluated at data values $\sqrt{Var(y_{k_0} \mid x_{-1})|_{x_{-1}=(\tilde{y}_{t-L},...,\tilde{y}_{-1})}}$ $t=0,...,n$, where y_t denotes data and y_{k0} denotes the kth element of the vector y_0, $k=1,...,M$. The volatility application involves estimating an unobserved state variable conditional on all past and present observables. Hence, filtering obtains [16], where y^* is the contemporaneous unobserved variable and x^* is the contemporaneous and lagged observed variables. Applications are portfolio optimization/minimization, option pricing and hedging.

3.1 Tail Probabilities, the Power Law and Extreme Values

The distribution of financial time series (returns) exhibits fatter tails than those of a normal distribution. The distribution for the latent volatility is more lognormal than normal. Hence, financial variables are four times more likely to experience big moves than the normal distribution would suggest. The power law, as an alternative to assuming normal distributions, asserts that it is approximately true that the value of a variable, υ, has the property that when y is large $\text{Prob}(\upsilon > y) = Ky^{-\alpha}$ where K and α are constants. A quick test is a plot of $\ln[\text{Prob}(\upsilon > y)]$ against $\ln y$. Evidence that the power law to hold is that this logarithm of the probability of the series changing more than y standard deviations is approximately linearly dependent on $\ln y$ for $y \geq 3$. Furthermore, the

[4] Filtered volatility is a data-dependent concept and the dynamic system must be sampled at the name frequency as the data to determine the density.

extreme value theory (EVT) estimates the tails of the volatility distributions [21]. EVT is a way of smoothing the tails of the probability distribution of daily changes. Value at Risk (VaR) and Expected Shortfall (ES) can be calculated and reflect the shape of the tail of the distribution. High confidence levels VaR and ES are available from EVT.

3.2 Volatility clustering

Volatility show clustering of periods of volatility, i.e. large (small) movements followed by further large (small) movements (shock persistence). In the financial literature, the lumpiness is called volatility clustering. Hence, a turbulent (tranquil) trading day (period) tends to be followed by another turbulent day (period). The implication is that volatility shocks today will influence the expectation of volatility for many periods in the future (shock persistence) and there are time varying return fluctuations in the markets.

3.3 Volatility exhibits persistence

The clustering of large and small movements (of either sign) from price movement processes is a well-documented feature in equity markets. To make a precise definition of volatility persistence let the expected value of the variance of returns *k* periods in the future be defined as $E_t\left(r_{t+k}-\mu_{t+k}\right)^2$ where *r* is the return and μ is the mean. The forecast of future volatility then depends upon information in today's information set such as today's return. Volatility is said to be persistent if today's return has a large effect on the forecast variance for many periods in the future. A measure of the persistence of volatility is the half-life. That is, the time it takes for the volatility to move half way back towards its unconditional mean following a deviation from it and can be expressed as $\tau = k : \left|h_{t+k|t} - \sigma^2\right| = \frac{1}{2}\left|h_{t+1|t} - \sigma^2\right|$. Alternatively, SV model volatility persistence can be studied by inspection of correlograms (*Q*-statistics) or the Breusch-Godfrey Lagrange multiplier test. Significant *Q*-statistics and χ^2 statistics suggest persistence.

3.4 Volatility is mean reverting

Mean reversion in volatility is generally interpreted as meaning that there is a normal level of volatility to which volatility will eventually return. In contrast, volatility clustering (persistence) implies that volatility comes and goes. Hence, mean reversion in volatility means that very long forecasts of volatility should all converge to the same normal level of volatility, no matter when they are made. The implicit interpretation is that mean reversion in volatility shows that current information has no effect on the long run forecast. Hence, periods of high volatility will eventually give way to more normal volatility, and similarly, periods of low volatility will be followed by a rise in volatility. More precisely, mean reversion implies that current information has no effect on the long run forecast. Hence, $p\lim_{k\to\infty}\theta_{t+k|t} = 0,\ \textit{for all } t$, and which is also expressed as $p\lim_{k\to\infty} h_{t+k|t} = \sigma^2 < \infty,\ \textit{for all } t$. Furthermore, note that option prices are generally viewed as consistent with mean reversion. That is, under simple assumptions of option pricing, the implied volatilities of long maturity options are less volatile than short maturity options (loser to long run average volatility).

3.5 Volatility asymmetry (leverage)

For equity market returns, it is plausible that positive and negative shocks have a different impact on volatility. This asymmetry is sometimes ascribed to a leverage effect and sometimes to a risk premium effect. For the leverage effect, as the price of a stock rises, its debt-to-equity ratio decreases, lowering the volatility of returns to equity holders. For the risk premium effect, news of increasing volatility reduces the demand for a stock because of general risk aversion among market participants. Hence, the stock value decline is normally followed by an increase in volatility as forecasted by news. Alternatively, price movements are negatively correlated with volatility suggesting that volatility increases (decreases) if the previous day returns are negative (positive) [6], [11]. Moreover, these authors also state that leverage effect happens because the fall (rise) in stock price causes leverage and the financial risk of the firm to increase (decrease).

3.6 Long Memory in Volatility

Financial time series exhibit long memory or persistence for volatility. Bailie et al. [3] states “The presence of long memory can be defined in terms of the persistence is consistent with an essentially stationary process, but where the autocorrelation takes far longer to decay than the exponential rate associated with the ARMA process”. The stochastic volatility (SV) models use long memory for modelling persistence. The autocorrelations for squared returns provide insights into the long memory characteristics of volatility measures. If the autocorrelations remain positive for very long lags, the long memory effect is present [22]. Moreover, explicit SV model volatility must exhibit the characteristics of long memory.

4 European Examples: FTSE100 index and Equinor asset

The daily analyses cover the period from the end of 2010 until November 2019, a total of 9 years and 110 consecutive months giving 2,325 returns for the two series. Price series are non-stationary and stationary logarithmic returns from all three series are therefore used in the analysis. Any signs of successful SV-model implementations for the markets indicate non-predictive market features and a minimum of weak-form market efficiency. Consequently, the markets are applicable for enhanced risk management activities.

4.1 Equity Summaries

Summary statistics for the two time-series are presented in Table 1. Both the FTSE100 spot index and the Equinor spot price series have small positive average returns (positive drift). The standard deviation for the index (portfolio) 0.928 is naturally lower than the single asset Equinor asset 1.587 (the index elements have a positive correlation less than 1), reporting lower risk. The maximum (3.9) and minimum (-6.2) numbers confirm lower risk for the FTSE100 index relative to the asset Equinor (a maximum of 8,7 and a minimum of -7.6). The FTSE index reports a negative skewness coefficient indicating that the return distributions are negatively skewed. In contrast, the asset Equinor reports a positive skewness suggesting a positively skewed distribution (more extreme positive price movements). The kurtosis coefficients are relatively high positive for both series (> 0), indicating a relatively peaked distributions with heavy tails. The FTSE100 series is peakier than the Equinor series suggesting that the FTSE100 index has more observations close to the unconditional mean. The JB normal test statistics [25] suggest non-normal return distributions. In contrast, the quantile normal test statistics suggest more normal distributed returns. Serial correlation in the mean equation is strong and the Ljung-Box *Q*-statistic [28] is significant for both series. Volatility clustering using the Ljung-Box test statistic for squared returns (Q^2) and ARCH statistics seems to be present. The ADF [13] and the Phillips-Person test statistics reject non-stationary series and the KPSS [27] statistic (12 lags) cannot reject stationary series. The RESET [31] test statistic, covering any departure from the assumptions of the maintained model, is not significant (stability). Finally, the BDS [7] test statistics report highly significant data dependence for all integrals (*m*). Figure 1 reports prices and returns and correlogram for the returns and squared/absolute returns. The correlogram for returns show only weak dependence while the correlogram for squared and absolute returns indicate substantial data dependence. The price change (log returns) data series (top), show that the level of volatility seems to change randomly but shows a time varying nature typically for financial markets.

FTSE100 Index	Mean (all)/ M (-drop)	Median Std.dev.	Maximum / Minimum	Moment Kurt/Skew	Quantile Kurt/Skew	Quantile Normal	Jarque - Bera	Serial dependence Q(12)	Q2(12)
	0.01117	0.04470	3.9429	3.4296	0.19367	4.1806	1081.5630	29.347	789.21
		0.91763	-6.1994	-0.35950	-0.04916	{0.1236}	{0.0000}	{0.0030}	{0.0000}
	BDS-Z-statistic (ε = 1)				KPSS	Ph-Perron	Augmented	ARCH	RESET
	m=2	m=3	m=4	m=5	I + Trend	I + Trend	DF-test	(12)	(12;6)
	10.0916	13.3062	15.2696	17.0384	0.02397	-45.09650	-44.6335	305.916	44.0178
	{0.0000}	{0.0000}	{0.0000}	{0.0000}	{0.7238}	{0.0000}	{0.0000}	{0.0000}	{0.0000}
Equinor Asset	Mean (all)/ M (-drop)	Median Std.dev.	Maximum / Minimum	Moment Kurt/Skew	Quantile Kurt/Skew	Quantile Normal	Jarque - Bera	Serial dependence Q(12)	Q^2(12)
	0.01240	-0.02732	8.6859	2.4184	0.24559	5.4668	523.7562	18.050	475.30
		1.58765	-7.6262	0.15990	0.01852	{0.0650}	{0.0000}	{0.1140}	{0.0000}
	BDS-Z-statistic (e = 1)				KPSS	Ph-Perron	Augmented	ARCH	RESET
	m=2	m=3	m=4	m=5	I + Trend	I + Trend	DF-test	(12)	(12;6)
	10.8927	12.5739	13.8733	14.9152	0.03720	-47.15972	-35.2619	208.747	7.790726
	{0.0000}	{0.0000}	{0.0000}	{0.0000}	{0.3645}	{0.0000}	{0.0000}	{0.0000}	{0.2538}

Table 1. Characteristics from the European Equity Markets

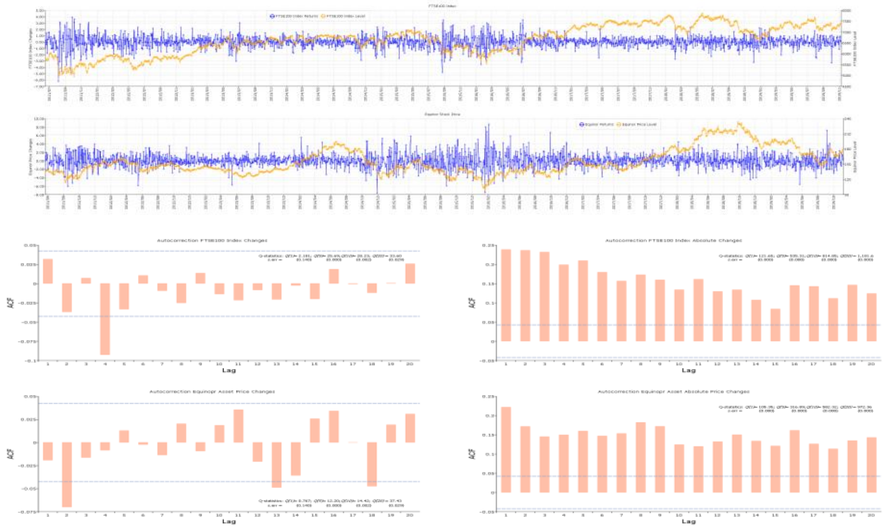

Fig. 1. FTSE100 Index (London) and Equinor Asset Prices (Oslo) for the period 2010 – 2019

4.2 The Stochastic volatility models for the European equities

The y_t is the percentage change (logarithmic) over a short time interval (day) of the price of a financial asset traded on an active financial market. The SV model implementation establishes a mapping between the statistical and the scientific models. The adjustment for actual number of observations and number of simulations is carefully logged for final model assessment. The SV model from equation 1 is estimated using efficient method of moments (EMM). The BIC [33] optimal SV model from parallel runs are reported in Table 2. The mode, mean and standard deviation are reported. For the two equity markets, a factor SV model produces acceptable model test statistics, reported at the bottom of Table 2. The objective function accuracy is -5.2 and -4.7 for the FTSE100 index and the Equinor asset, respectively, with associated χ^2 test statistics of 0.26 (4 *df*) and 0.20 (3 *df*). The MCMC log-posterior are reported in Figure 2. The model does not fail the test of over identified restrictions at the level of 10%, the chains are choppy, and the densities are close to normal, all factors suggesting that the SV model is appropriate for the two equity markets. The long-simulated realization of the state vector, as a-by product of the estimated SV model, establishes a functional form of the conditional distribution. The SNP methodology obtains a convenient representation of one-step ahead conditional variance $\hat{\sigma}_t^2$ *of* $\hat{y}_{t+1}$ *given* $\{\hat{y}_\tau\}_{\tau=1}^{t}$. Running regressions for V_i on $\hat{\sigma}_t^2$, $\hat{y}_t$ *and* $|\hat{y}_\tau|$ and a generous number of lags of theses series, we obtain calibrated functions that give predicted values of $V_{it} | \{y_\tau\}_{\tau=1}^{t}$, $t = 1,2$ on the observed data series. Figure 3, reporting the last 60 days in 2019, shows the two latent volatility factors for the observed data points. The plots indicate that V_1 is slowly moving while V_2 is moving considerably faster. It is quite clear that the slowly persistent factor V_1, leads the re-projected yearly volatility for both series. Figure 3 also reports the ordinary least square number for R^2 for FTSE100 index (Equinor asset) at a level of V_i, where i=1,2

FTSE100 Index Scientific Model				Equinor Asset Price Scientific Model			
Parameter values Scientific Model.			Standard	Parameter values Scientific Model.			Standard
θ	Mode	Mean	error	θ	Mode	Mean	error
$a0$	0.017578	0.022952	0.013395	$a0$	-0.023438	-0.016096	0.035251
$a1$	0.011719	0.002120	0.021067	$a1$	-0.054688	-0.062437	0.023304
$b0$	-0.324220	-0.344770	0.056698	$b0$	0.484380	0.308630	0.199130
$b1$	0.939450	0.934230	0.015442	$b1$	0.828120	0.802790	0.088097
$c1$	0	0	0	$c1$	0	0	0
$s1$	0.134770	0.129250	0.015180	$s1$	0.179690	0.167230	0.039989
$s2$	0.136720	0.130550	0.041149	$s2$	0.148440	0.139710	0.059609
$r1$	-0.785160	-0.731250	0.066906	$r1$	-0.531250	-0.377280	0.203640
$r2$	0.511720	0.446280	0.139440	$r2$	0.656250	0.513340	0.256090
Distributed Chi-square (no. of freed			$\chi 2(4)$	Distributed Chi-square (no. of freed			$\chi 2(3)$
Posterior at the mode			-5.2477	Posterior at the mode			-4.6865.
Chi-square test statistic			{0.2628}	Chi-square test statistic			{0.1962}

Table 2. Scientific Stochastic Volatility Characteristics: θ-parameters

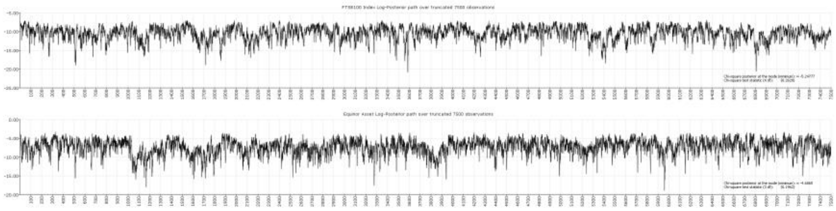

Fig 2. MCMC posterior chain from 250 *k* optimal SV model (R = 75.000)

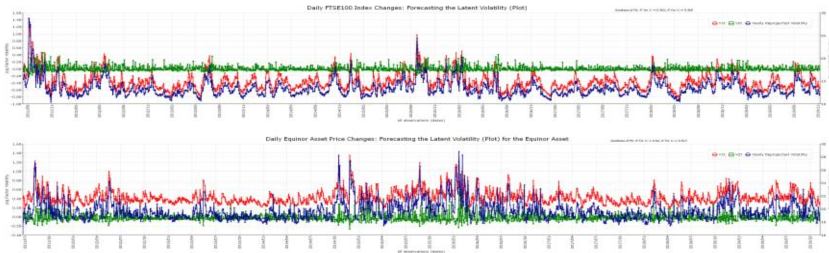

Fig. 3. Conditional Volatility from Observables and Kalman Filtered Volatility (daily)

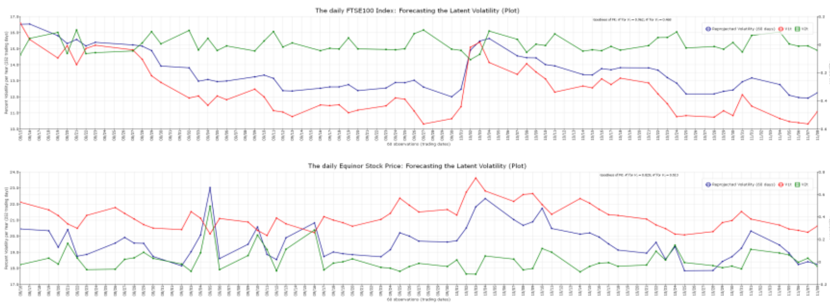

Fig. 4. FTSE100 index (top) and Equinor Stock (bottom) Factor Volatility Paths (last 60 days)

of 96.2% (82%) and 46.8% (51.3%), respectively. Obviously, the slowly moving V_1 factor, showing persistence, is the main contributor to yearly volatility. V_2 moves much faster showing strong mean reversion, absorbing shocks.

4.3 Volatility characteristics for the European Equities

The volatility factors in Figure 3 and 4 seem to model two different flows of information to the equity markets. One slowly mean reverting factor provides volatility persistence and one rapidly mean reverting factor provides for the tails [8]. The factor for the FTSE100 index is clearly moving slower than for the Equinor asset. In contrast to the crash of 1987 which was attributed to a large realization of the mean reverting factor V_2, the period 2011 to 2019 does not show large realization of V_2, but rather much more to the slowly moving factor V_1. In accordance with the plots, the period from 2011 to 2019 seems to show slow and persistent changes to volatility. However, for the Equinor asset oil shocks have shown some major contributions to volatility. For example, the shock in May 2019 is only temporary and the volatility from the shock, show strong mean reversion (V_2).

Comparing Figure 1 and 3, the two synchronous plots show that when returns become wider (narrower) volatility increases (decreases). Moreover, turbulent (wide returns) days tend to be followed by other turbulent days, while tranquil (narrow returns) tend to follow other tranquil days (clustering). As should be expected, the volatility is clearly higher for the Equinor asset than for the FTSE100 index. Furthermore, the volatility seems to increase more from negative returns than from positive returns. Volatility densities for the FTSE100 index and the Equinor asset series suggest lognormal densities. As suggested above, the density for Equinor shows both narrower and higher volatility density than the FTSE100 index. Furthermore, the power law $\left(\text{Prob}(\upsilon > x) = Kx^{-\alpha}\right)$ providing an alternative to the normal distributions, seems approximately true for the volatility. Finally, Figure 5 reports the correlogram for the FTSE100 index and the Equinor asset. The correlograms indicates substantial dependence suggesting both clustering and persistence as well as making volatility predictions more relevant.

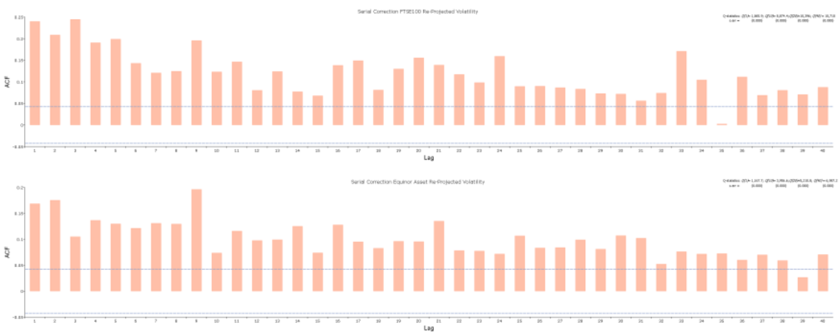

Fig. 5. Conditional Volatility from Observables and Kalman Filtered Volatility (daily)

Tail properties, the Power law and Extreme values. The power law, an alternative to assuming normal distributions, is applied to the reprojected volatility $\left(e^{(V_1+V_2)}\right)$ for the FTSE100 Index and Equinor asset. The power law asserts that, for many variables, it is approximately true that the value of the variable, □, has the property that when x is large $\text{Prob}(\upsilon > x) = Kx^{-\alpha}$ where K and α are constants. The relationship implies that $\ln\left[Prob(\upsilon > x)\right] = \ln K - \alpha \ln x$, and a test of whether it holds by plotting $\ln\left[Prob(\upsilon > x)\right]$ against $\ln x$. The values for *ln(x)* and *ln[Prob(v > x)]* for the FTSE100 index and the Equinor asset show that the logarithm of the probability of a change by more than *x* standard deviations is approximately linearly dependent in *ln(x)* for $x \geq 3$. Hence, for both the FTSE100 index and the Equinor asset the power law holds for the re-projected volatility. Regressions show the estimates of K and α are as follows: for FTSE100 (Equinor) $K = e^{-2.274}$ and $\alpha = 2.147$ ($K = e^{-0.379}$ and $\alpha = 3.369$). A probability estimate of a volatility greater than 3 (6) standard deviations is $0.103 \cdot 3^{-2.147} = 0.0097(0.97\%)$ $\left(0.103 \cdot 6^{-2.147} = 0.0022(0.22\%)\right)$ and $0.685 \cdot 3^{-3.369} = 0.0169(1.69\%)$ $\left(0.685 \cdot 6^{-3.5369} = 0.0016(0.16\%)\right)$ for the FTSE100 index and the Equinor asset, respectively. The extreme value theory takes us a bit further. Setting the *u* to the 90 percentiles of the filtered volatility series of FTSE100 (*u*=15.55) and Equinor *(u*=21.65). The FTSE100 index reports optimal β = 1.648 and ζ = 0.119 with an associated maximum value for the log-likelihood function of -341.6. The Equinor series reports optimal β = 1.3067 and ζ = 0.0514 with an associated maximum value for the log-likelihood function of -278.3. The probability that the FTSE100 index re-projected volatility will be greater than 20 (30) is 0.9634% (0.025%). The VaR with 99% (99.9%) confidence limit is 19.92 (25.67). Hence, the 99.9% VaR estimate is about 0.892 times lower than the highest historic re-projected volatility. The 99% (99.9%) expected shortfall (ES) estimate is 22.38 (28.92). Furthermore, for the FTSE100 index, the unconditional probability for a volatility greater than 15.5356 (*u*) is equal to 0.46%. Similarly, the probability that the Equinor asset re-projected volatility will be greater than 20 (30) is 37.07% (0.9%). The VaR with 99% (99.9%) confidence limit is 24.85 (28.45). Hence, the 99.9% VaR estimate is about 1.005 times higher than the highest historic filtered volatility for the Equinor asset. The 99% (99.9%) ES estimate is 26.265 (30.068). Finally, for the Equinor asset, the unconditional probability for volatility greater than 21.798 (*u*) is equal to 0.68%. As Var and ES are attempts to provide a single number that summarizes the volatility tails giving the market participants an indication of the risk to which they are exposed. The FTSE100 index shows that a daily volatility greater than 20 is only 0.9634% while the Equinor asset, as a single asset, shows that a daily volatility greater than 20 is 37.06%. Hence, EVT and the power law, reporting VaR and ES values, summarises tail properties that indicate the risk for the market participants. For market participants, inverting the unconditional probability for volatility and setting a 1% limit for the change of unconditional probability, will list associated investments alternatives.

Volatility clustering. The BDS independence test statistic [7] is a portmanteau test for time-based independence in a series. The probability of the distance between a pair of points being less or equal to epsilon (ε) should be constant ($c_m(\varepsilon)$). The BDS test statistics, where ε is one standard deviation and the number of dimensions is 10, reports that for both the FTSE100 index and the Equinor asset, the data strongly rejects the hypothesis that the observations are independent. The FTSE100 index shows a higher BDS dependence than the Equinor asset. Moreover, the SV model reports volatility serial correlation with the *SV* coefficient b_1. The correlation is much stronger for the FTSE100 index (b_1 =0.94) then for the Equinor asset (b_1 =0.83). The $b_1 > 0.8$ will accommodate volatility clustering. This is also visible in the above Figure 3 showing longer periods of high/low volatility for the FTSE100 index than for the Equinor asset (choppier).

Persistence in Volatility. Figure 4 reports the autocorrelation and partial autocorrelation functions up to 20 lags for the FTSE100 Index and the Equinor asset. The pattern of temporal dependence is different for the two volatility factors, V_1 and V_2. V_1 shows strong temporal dependence while V_2 shows close to zero temporal dependence. The re-projected volatility $\left(e^{(V_1+V_2)}\right)$ has inherited the temporal dependence from V_1, suggesting strong persistence in volatility. The correlograms show that FTSE100 index show higher correlation for the first lags, 0.940 versus 0.824 for the Equinor asset. However, from lag nine and higher the Equinor asset show higher serial correlations. Running the Breusch-Godfrey serial correlation LM test (Godfrey, 1988) also report strong serial correlation up to lag 20 of 1869.76 (χ^2(20)={0.000}) and 1399.42 (χ^2(20)={0.000}) for the FTSE100 index and the Equinor asset, respectively. Hence, the re-projected volatility for both the FTSE100 index and the Equinor asset, show strong volatility persistence.

Volatility is mean reverting. A battery of unit root tests together with a variance ratio test (martingales) are used to test for mean reversion for the re-projected volatility. For example, the FTSE100 index (Equinor) report an ADF statistic of -9.4 (-7.7). Hence, the ADF statistics report significant mean reversion at the 1% level. Furthermore, all unit-root test statistics suggest stationary and mean reverting series. The overlapping variance ratio test [29], examines the predictability of time series data by comparing variances of differences in the data (returns) calculated over different intervals. If we assume the data follow a random walk, the variance of a period difference should be times the variance of the one-period difference. The FTSE100 index (4.399) and the Equinor asset (5.588) both reject that the volatility is a martingale, suggesting mean reversion.

Asymmetry in Volatility. The asymmetry and the leverage effects are the negative correlation between the shocks of return and the subsequent shocks on volatility. Hence, after a negative return shock, we expect volatility to increase while after a positive shock on returns we should observe a decrease in volatility. Studying the volatility changes following return shocks gives some information regarding this proposition. Dividing the volatility from positive and negative returns show for the FTSE100 index (Equinor asset) an average increase in volatility from negative shocks of 2.057 (1.912) and from positive shocks of -1.875 (-1.897). Hence, negative return shocks increase average volatility while positive return shocks decrease average volatility. To statistically test for the change in volatility from negative and positive returns, we run an OLS regression on the change in daily volatility on returns and lagged returns. For the FTSE100 index (Equinor asset) the regression reports a coefficient from the returns equal to -2.206 (-0.071) and -2.823 (-1.671) for lagged returns, all significant at the 5% level. That is, the two series show that negative returns seem to increase volatility while positive returns seem to reduce volatility. Furthermore, the correlation coefficients between returns and synchronous (and lagged) re-projected volatility is -0.541 (-0.683), and -0.0164 (-0.6831) for the FTSE100 index and the Equinor asset, respectively, suggesting negative return asymmetry for both series.

Long memory. Long memory is associated with both clustering and persistence. By using fractional differencing with traditional ARMA specifications, the ARFIMA model allows for flexible dynamic patterns for the re-projected volatility. For the FTSE100 index, the ARFIMA (2,*d*,0) model specification estimate $d = 0.3043$ suggests slow autocorrelations and partial autocorrelations decay (hyperbolically). For the Equinor asset, the ARFIMA (2,*d*,0) model specification estimate $d = 0.3571$ suggests the same slow autocorrelations and partial autocorrelations decay. The ARFIMA model therefore specifies two slowly decaying series with long-run dependence (long memory).

4.4 Step Ahead Volatility Predictions for European Equities

The SNP methodology obtains a convenient representation of one-step ahead conditional variance $\hat{\sigma}_t^2$ *of* $\hat{y}_{t+1}$ *given* $\{\hat{y}_\tau\}_{\tau=1}^{t}$. Running regressions for V_{it} on $\hat{\sigma}_t^2$, $\hat{y}_t$ *and* $|\hat{y}_\tau|$ and a generous number of lags of theses series, we obtain calibrated functions that give step ahead predicted values of $V_{it} \,|\, \{y_\tau\}_{\tau=1}^{t},\ t = 1,2$ at the data points. A static forecast for the FTSE100 index and the Equinor asset is done in Figure 6. The estimation period is from 2010 to January 1st, 2019 and the static forecasting period from January 1st, 2019 to November 8th 2019. For a “good” measure of fit, using the Theil inequality coefficient (bias, **v**ariance and covariance portions) the bias and variance should be small so that most of the bias is concentrated on the covariance proportion. The covariance proportion for re-projected volatility for the FTSE100 index (Equinor asset) is 0.966 (0.918). For the main contributor to re-projected volatility for both series, factor V_1, the covariance portion of the Theil inequality coefficient is even higher.

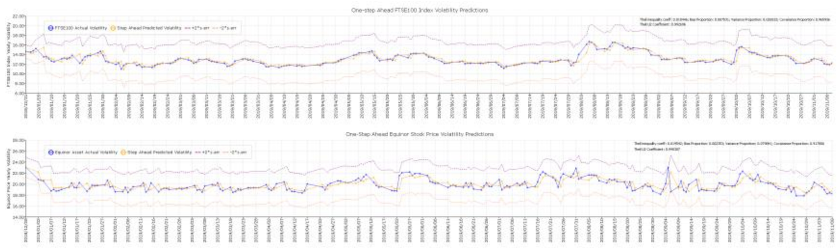

Fig. 6. Static Forecasts for the FTSE100 Index and the Equinor Asset 2019.

5 Summary and Conclusions

The main objective of this paper has been to characterize a good volatility model by its ability to forecast and capture the commonly held stylized facts about equity market volatility. The stylized facts include such things as heavy tails, persistence, mean reversion, asymmetry (negative return innovations suggest higher volatility), and long memory. The characteristics indicate substantial data dependence in the volatility. The paper shows that the re-projected volatility contains all these characteristics and that this data dependence suggests an ability for volatility predictions to enhance risk management, portfolio timing and selection, market making and derivative pricing for speculation and hedging in equity markets.

The paper has used the Bayesian M-H estimator and a stochastic volatility representation for European financial equity markets. The methodology is based on the simple rule: compute the conditional distribution of unobserved variables given observed data. The observables are the asset prices and the un-observables are a parameter vector, and latent variables. The inference problem is solved by the posterior distribution. Based on the Hammersley-Clifford [23] theorem, $p(\theta,x|y)$ is completely characterized by $p(\theta|x,y)$ and $p(x|\theta,y)$. The distribution $p(\theta|x,y)$ is the posterior distribution of the parameters, conditional on the observed data and the latent variables. Similarly, the distribution $p(x|\theta,y)$ is the smoothing distribution of the latent variables given the parameters. The MCMC approach therefore extends model findings relative to non-linear optimizers by breaking the "curse of dimensionality" by transforming a higher dimensional problem, sampling from $p(\theta_1,\theta_2)$, into easier problems, sampling from $p(\theta_1|\theta_2)$ and $p(\theta_2|\theta_1)$ (using the Besag [5] formula).

Although price processes are hardly predictable, the variance of the forecast error is clearly time dependent and can be estimated by means of observed past variations. The results suggest that volatility can be forecast. The stochastic volatility models are therefore an area in empirical financial data modelling that is fruitful as a practical descriptive and forecasting device for all participants/managers in the financial services sector, together with a special emphasis on risk management (forecasting/ re-projections and VaR/ES), portfolio management and derivative innovations. Irrespective of markets and contracts, Monte Carlo Simulations should lead us to more insight into the nature of the price processes describable from stochastic volatility models. Finally, static predictions of the re-projected volatility suggest a relatively good fit.

References

1. Andersen, T.G., Stochastic autoregressive volatility: a framework for volatility modelling, Mathematical Finance, 4, pp. 75-102 (1994)
2. Andersen, T.G., L. Benzoni, and J. Lund, Towards an empirical foundation for continuous-time models, Journal of Finance, 57, pp. 1239-1284 (2002)

3. Bailie, R.T., T. Bollerslev and H.O. Mikkelsen, Fractionally integrated generalized auto-regressive conditional heteroskedasticity, Journal of Econometrics (74), pp. 3-30 (1996)
4. Bau III, David and Lloyd N Trefethen, Numerical Linear Algebra, Philadelphia: Society of Industrial and Applied Mathemathics (1997)
5. Besag, J., Spatial interaction and the statistical analysis of lattice systems (with discussion), Journal of Royal Statistical Society Series B 36, pp. 192-326 (1974)
6. Black, F., Studies of stock market volatility changes. Proc. 1976 Meetings of the American Statistical Association, Business and Economics Statitics Section, pp. 307-327 (1976)
7. Brock, W. A., W.D. Dechert, J.A. Scheinkman, and B. LeBaron, A test for independence based on the correlation dimension, Econometric Reviews, 15,197-235 (1996)
8. Chernov, M., A.R. Gallant, E. Ghysel, and G. Tauchen, Alternative models for stock price dynamics, Journal of Econometrics, 56, PP. 225-257 (2003)
9. Chernozhukov, Victor, and Han Hong, An MCMC Approach to Classical Estimation, Journal of Econometrics 115, pp. 293-346 (2003)
10. Clark, P. K., A subordinated stochastic Process model with finite variance for speculative prices, Econometrica, 41,135-156 (1973)
11. Christie, A. A., The stochastic behaviour of common stock variances: value, leverage and interest rate effects. Journal of Financial Economics (10), pp. 407-432 (1982)
12. Dickey. D.A. and W.A. Fuller, Distribution of the estimators for autoregressive time series with a unit root, Journal of the American Statistical Society (74),(366), pp. 427-431 (1979)
13. Durham, G., Likelihood-based specification analysis of continuous-time models of the short-term interest rate, Journal of Financial Economics, 70, pp. 463-487 (2003)
14. Gallant, A.R. and G. Tauchen, A Nonparametric Approach to Nonlinear Time Series Analysis., Estimation and Simulation, in David Brillinger, Peter Caines. John Geweke, Emanuel Parzan, Murray Rosenblatt, and Murad S.Taqqu eds. New Directions in Time Series Analysis, Part II. New York: Springer-Verlag, pp. 71-92 (1992)
15. Gallant, A.R., D.A. Hsieh and G. Tauchen, Estimation of stochastic volatility models with diagnostics, Journal of Econometrics, 81, pp. 159-192 (1997)
16. Gallant, A.R. and J.R. Long, Estimating stochastic differential equations efficiently by minimum chi-squared, Biometrika, 84, pp. 125-141 (1997)
17. Gallant, A. R. & R.E. McCulloch, GSM: A Program for Determining General Scientific Models, Duk, 84, University (http://econ.duke.edu/webfiles/arg/gsm) (2011)
18. Gallant, A.R. and G. Tauchen, Reprojecting Partially Observed Systems with Application to Interest Rate Diffusions, Journal of the American Statistical Association, 93 (441), pp. 10-24 (1998)
19. Gallant, A.R. and G. Tauchen, Simulated Score Methods and Indirect Inference for Continuous Time Models, in Yacine Aït-Sahalia, Lars Petter Hansen, eds. Handbook of Financial Econometrics, North Holland, Chapter 8, pp. 199-240 (2010)
20. Gallant, A.R. and G. Tauchen, EMM: A Program for Efficient Methods of Moments Estimation, Duke University (http://econ.duke.edu/webfiles/arg/emm) (2010)
21. Gnedenko, D.V., Sur la Distribution limité du terme d'une série aléatoire, Annals of Mathematics (44), pp. 423-453 (1943)
22. Granger C. and Z. Ding, Some Properties of Absolute Returns. An Alternative Measure of Risk, Annals of Economics and Statistics (40), pp. 67-91 (1995)
23. Hammersley, J. and P. Clifford, Markov fields on finite graphs and lattices, Unpublished manuscript (1970)
24. Hansen, L.P., Large Sample Properties o Generalized Method of Moments Estimators, Econometrics 50, pp. 1029-1054 (1982)

25. Jarque, J.B. and A-K. Bera, Efficient tests for normality, homoscedasticity and serial independence of regression residuals, Economic Letters, 6, (3), pp. 255–259 (1980)
26. Johnson, N.L., S. Kotz and N. Balakrishnan, Continuous Univariate Distributions, in Wiley & Sons, 1995 – 752 (1970)
27. Kwiatowski, D., P.C.B. Phillips, P. Schmid and T. Shin, Testing the null hypothesis of stationary against the alternative of a unit root: How sure are we that economic series have a unit root, Journal of Econometrics, Vol. 54, pp 159-178 (1992)
28. Ljung, G.M. and G.E.P. Box, On a Measure of Lack of Fit in Time Series Models, Biometrika, 65, 297-303 (1978)
29. Lo, A.W. and C. MacKinlay, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, The Review of Financial Studies, 1 (1), pp. 41–66 (1988)
30. Musiela, M and M. Rutkowski, Martingale Methods in in Financial Modelling, Springer-Verlag Berlin (1997)
31. Ramsey, J. B., "Tests for Specification Errors in Classical Linear Least Squares Regression Analysis". Journal of the Royal Statistical Society, Series B. 31 (2), pp. 350–371 (1969)
32. Rosenberg, B., 1972, The behavior of random variables with nonstationary variance and the distribution of security prices, Unpublished paper, Research Program in Finance, University of California, Berkeley
33. Schwarz, G., Estimating the Dimension of a Model, Annals of Statistics, 6, pp. 461-464 (1978)
34. Shepard, N., Stochastic Volatility: Selected Readings, Oxford University Press (2004)
35. Tauchen, G. and M. Pitts, The price variability volume relationship on speculative markets, Econometrica, pp. 485-505 (1983)
36. Taylor, S., Financial returns modelled by the product of two stochastic processes – a study of daily sugar prices 1961-79. In Anderson. O. D. (ed.) Time Series Analysis: Theory and Practice, 1, pp. 203-226, Amsterdam, North-Holland (1982)
37. Taylor, Stephen, Asset Price Dynamics, Volatility, and Prediction, Princeton University Press (2005)