Quantitative and Qualitative Analysis in Social Sciences
Volume 1, Issue 2, 2007, 33-54
ISSN: 1752-8925

An Analysis of Oil Price Behaviour

Nigel Meade[a] **Anna Cooper**
Tanaka Business School, Imperial College

Abstract

After a brief summary of oil price behaviour from 1946 onwards, the problem of modelling daily oil prices from 1990 onwards is addressed. The arbitrage pricing continuous time models from the literature are compared with discrete time ARMA models. The evidence for the mean reversion process favoured by continuous time models is found to be reduced if the time series model takes GARCH effects and uses a non-Gaussian error term with fatter tails.

JEL Classifications: C32, C53.
Keywords: Time series analysis, oil prices, density forecasting, interval forecasting.

[a] Tanaka Business School, Imperial College, Exhibition Road, London SW7 2AZ; tel.: +44 (0)20 7594 9116, email: n.meade@imperial.ac.uk

1 Introduction

The intention of this paper is to compare the models of oil price behaviour proposed in the literature. Additional models based on data analysis will also be considered. The motivation for model building is the increasing use of oil based derivatives, such as futures, for risk management in the oil market and for investment purposes by other investors.

One modelling philosophy is dictated by arbitrage pricing models where the motivation is to provide a pricing framework for a derivative or future. These continuous time models include geometric Brownian motion and mean reversion models. An alternative philosophy is data driven, where a model is chosen from a universe of models such as the ARIMA framework according to a goodness of fit procedure.

2 Oil Price History and Data

Oil price data is easily available from 1946 onwards; however, between 1946 and 1970 the oil price crept up from 1.63 US$ per barrel ($/bl) to 3.39$/bl. The history from 1970 onwards is summarised in Figure 1.

Figure 1

An annotated plot of the crude oil price from 1970 to 2006.

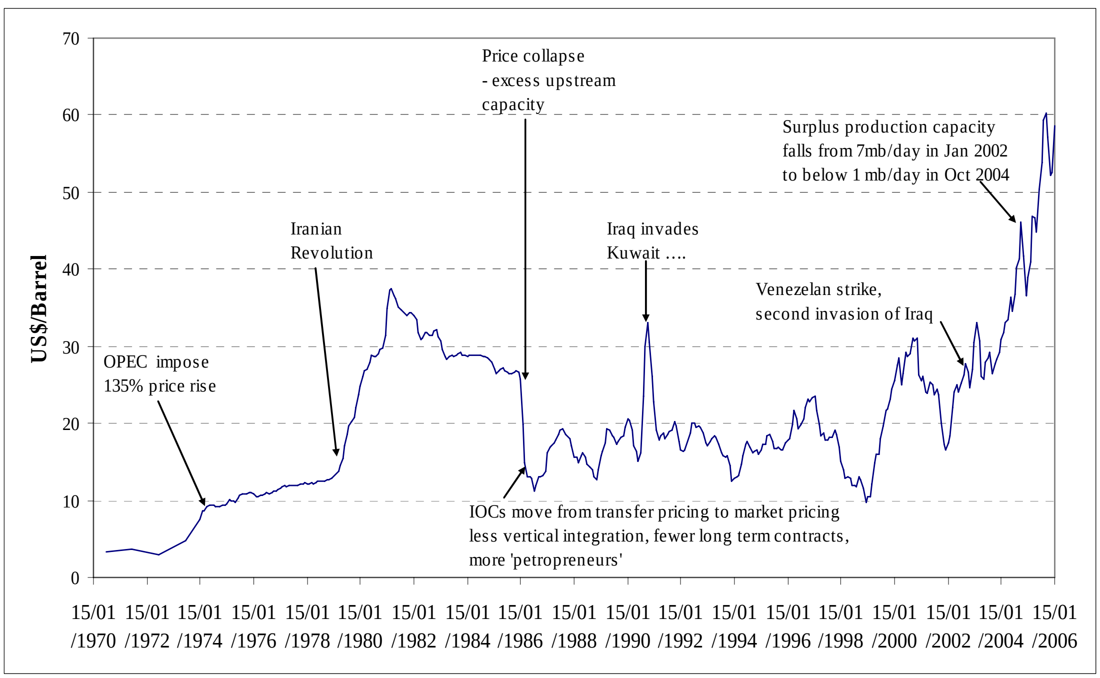

Various Middle Eastern crises led to the price exceeding 10$/bl in 1974. After reaching a peak of 37.50$/bl in March 1981, the prices fell back to the 10-20$/bl band where it stayed until 2000. Stevens (2005) attributes the sharp price drop in 1986 to excess upstream capacity. In the late 1980s, the international oil companies (IOCs) began to use long-term contracts less and to use the market more. The markets became more important with the quoted prices of marker crude oils such as Brent, West Texas Intermediate (WTI) and Arab Light being used as a basis to price other crude oils. The period 2002 onwards shows an upward movement in prices - during this period surplus production capacity decreases from 7 million bbl per day (Jan. 2002) to less than one million bbl per day (Oct. 2004).

For our detailed analysis of the market, we use prices from April 1991 onwards. During this period the market is liquid and the price spike caused by the Iraqi invasion of Kuwait is avoided.

The precise data set consists of daily prices from 8 April 1991 to 5 April 2006 for Brent in US$/bl obtained from Datastream. The price data is plotted in Figure 2.

Figure 2

Daily data (April 1991 to April 2006) for prices, returns and volatility.

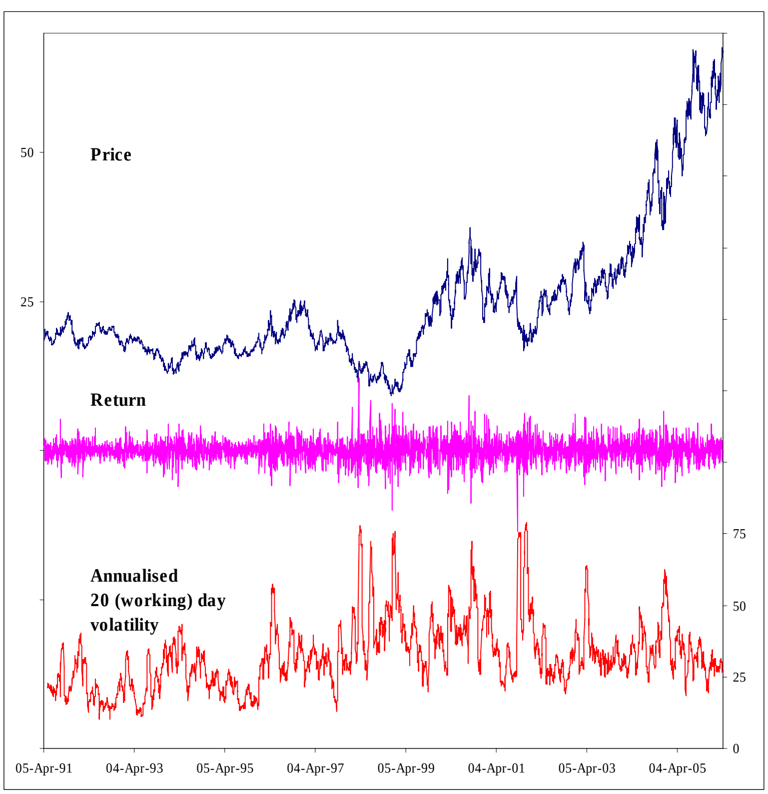

Daily returns are also shown, the return plot helps highlight the periods of high volatility, and these are further illustrated by the twenty (working) day volatility plot shown at the bottom of the graph.

3 Oil Pricing Models

3.1 Arbitrage Pricing Models

Brennan and Schwartz (1995) proposed one of the early oil pricing models as part of their suggestion for pricing natural resource investments. In effect, they were pricing a real option. They proposed using a geometric Brownian motion model for the price *S*:

$$\frac{dS}{S} = \mu \, dt + \sigma \, dz_t , \tag{1}$$

where dz_t is the increment in a Gauss-Wiener process with drift μ and instantaneous standard deviation σ.

Schwartz (1997) proposed three models with one, two and three factors. The factors are spot price, convenience yield and interest rate. The purpose of these models was to price futures, whereas here our pre-occupation is with the spot price (although this is the price of the futures contract closest to maturity). Thus we only need to consider the first of the three models proposed as the last two models collapse to (1) if we assume a constant convenience yield. The first model is a mean reversion (MR) model:

$$\frac{dS}{S} = \kappa\left(\mu - \ln\left(S\right)\right) dt + \sigma \, dz_t , \tag{2}$$

where the log of the oil price reverts to a long-term mean μ at a rate defined by κ.

Schwartz and Smith (2000) propose a combined long-term and short-term model; the log price is the sum of these components:

$$\ln\left(S_t\right) = \chi_t + \xi_t . \tag{3}$$

The short-term process is a reversion process to a zero mean:

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi \, dz_\chi , \tag{4}$$

and the long-term process is a geometric Brownian motion (in price):

$$d\xi_t = \mu_\xi dt + \sigma_\xi \, dz_\xi . \tag{5}$$

The two processes are correlated as follows:

$$dz_{\chi}\ dz_{\xi} = \rho_{\chi\xi} dt\ . \tag{6}$$

The models described in (1), (2), and (3)-(6) are the three continuous time models to be found in the literature (to the best of our knowledge). Other authors, for example Cortazar and Naranjo (2006), have proposed multi-factor models but they have not introduced another spot price process.

3.2 Discrete Time Models

Here we denote the return $Z_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$ - note that, as in general in financial econometrics, it is taken as given that Z_t is I(0) and S_t is I(1). Panas and Nini (2000) model oil product prices, the most general form of their model is an AR(1)-GARCH(1,1) in mean model. The AR(1) process (with GARCH in mean term) is

$$Z_t = \mu + \phi_1 Z_{t-1} + \gamma\sigma_{t-1}^{\lambda} + \varepsilon_t\ , \tag{7}$$

and the GARCH (1,1) process is

$$\sigma_t^2 = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \beta_1\sigma_{t-1}^2\ , \tag{8}$$

where σ_t^2 is the conditional variance of Z_t, $\alpha_0 > 0$, $\alpha_1 > 0$, $\beta_1 \geq 0$, and λ is typically 1 or 2. In their study of fourteen products on the Rotterdam and Mediterranean markets, they found six series with significant autoregressive terms ($\phi_1 \neq 0$) and a non-overlapping six series displaying a significant GARCH in mean effect ($\gamma \neq 0$). The GARCH effect was significant in all fourteen series.

Cabedo and Moya (2003) used ARMA models with the intention of using the model to predict value at risk (VaR). They fitted an ARMA(1,1) model to daily returns on Brent crude for 1992 to 1998. The fitted model is

$$Z_t = 0.016 + 0.99Z_{t-1} + \varepsilon_t + 0.94\varepsilon_{t-1}\ . \tag{9}$$

They found that this model gave better out of sample VaR results than the AR(1)-GARCH(1,1) model.

3.3 Discrete Time Models Derived from Continuous Prototypes

The Euler-Maruyama method (see Higham, 2001) is a well recognised means of converting continuous models to a discrete time formulation. The Brownian motion of (1) becomes

$$Z_t = \mu + \varepsilon_t , \tag{10}$$

where $\varepsilon_t \sim N\left(0, \sigma^2\right)$. The mean reversion model in (2) becomes

$$Z_t = \kappa\left(\mu - \ln\left(S_{t-1}\right)\right) + \varepsilon_t . \tag{11}$$

If we assume a single source of error model, the short- and long-term processes combined in (3) simplify to the following equation in discrete time:

$$Z_t = \kappa\left(\mu - \ln\left(S_{t-1}\right)\right) + \varepsilon_t + \theta\varepsilon_{t-1} . \tag{12}$$

3.4 Evidence against Gaussianity and Alternative Density Functions

Most of the literature refers to equity returns rather than commodity returns, so we will draw on this literature to suggest other possible models for oil returns. Among others, Osborne(1959) postulated that the behaviour of returns on equity was consistent with Brownian motion, and that asset returns followed a Gaussian distribution. However, empirical observation of return processes provides much evidence to refute the hypothesis of Gaussianity. Based on observed leptokurtosis in asset return data, Mandlebrot (1963) argued that a stable distribution was a better model than the Gaussian. Akgiray and Booth (1988) considered a stable-law model for individual US equity returns. They concluded that empirical tails were thinner than implied by the fitted stable distribution; a skewed distribution with fatter tails than the Gaussian was to be preferred.

The detection of persistent skewness in returns is dependent on the sampling properties of the skewness coefficient. Although Singleton and Wingender (1986) found evidence of skewness in monthly returns, they found little evidence of its persistence at the single equity or the portfolio level. Using bootstrapping to estimate the sampling distribution of the skewness estimate on the same data set, Muralidhar (1993) found that skewness did persist in a large proportion of equities but not at a portfolio level. Peiro (1999) corroborated the second part of this conclusion, finding no evidence of skewness persistence in a study of daily returns on stock market indices (equivalent to portfolios). In contrast, Mills (1995) finds both extreme kurtosis and positive skewness in the daily returns on three London FTSE indices (post 1988). The behaviour of these index returns was characterised by a thinner left than right tail.

Several probability density functions have been proposed as alternatives to the Gaussian. Since this study will consider density functions for conditional and unconditional distributions of returns, ease of parameter estimation has to be considered. This means that the stable law will not be considered here. To make for easier reading, the density functions for all the random variables considered are given in the Appendix.

1. **The Gaussian random variable** is included as a benchmark. Since the densities are estimated given a mean and variance, no parameters are estimated for the Gaussian.
2. **A mixture of Gaussian random variables**. Kon (1984) proposed this random variable as a model of stock returns. It can represent leptokurtosis and skewness if the means of the components are not identical. Using daily returns from thirty US stocks, Kon (1984) found that the log-likelihood of the mixture was greater than for the Student's *t*. Mixtures of two, three, and four Gaussians were used. In applying VaR to foreign exchange rates, Venkataraman (1997) demonstrated that a mixture of two Gaussians yielded a number of violations (observations falling into a 5% tail, i.e., having a *p*-value $\alpha = 0.95$) that was consistent with the value of α, using a likelihood ratio test. The Gaussian distribution generated significantly more violations. In general, *m* Gaussians can be mixed and the choice of *m* is decided by a likelihood ratio test. The number of estimated parameters is *3m-1* with two constraints to give the required mean and variance.
3. **A mixture of Gaussian and Laplace random variables**. The motivation behind the use of the Laplace distribution is to use its thicker than Gaussian tails to model empirical leptokurtosis. There are five estimated parameters subject to two mean and variance constraints.
4. **The Normal Inverse Gaussian random variable (NIG)**. The density function of this random variable has four parameters. Given the mean and variance of the data, this leaves two parameters free to describe the shape of the distribution. Barndorff-Nielsen (1997) proposed it as a model of stock returns. Rydberg (1999) used this density function for daily

returns of two US equities, Coca-Cola and Ford. Its use in risk analysis was demonstrated by Lillestøl (2000), who fitted the density function to returns on the S&P 500 and FT Actuaries indices. (Other transformations of the Gaussian include Tukey's *g* and *h* distributions, used by Mills (1995) to model daily returns on three London FTSE indices. Edgeworth-Sargan distributions are used by Mauleon and Perote (2000) to model stock market indices; these distributions use polynomials of Gaussian random variables.)

5. **The Skewed Generalised Student's *t* random variable (SGST)**. Theodossiou (1998) introduced this density function to model the empirical behaviour of financial time series. It is a skewed version of the Generalised *t* of McDonald and Newey (1988). Using daily data, Theodossiou (1998) found that the skewness was significant for two foreign exchange rates but not for the stock market indices examined. This distribution has several density functions nested within it. These include the Gaussian; the generalised error distribution (or the power exponential) which includes the Laplace; the Student's *t*; the Cauchy and the uniform. Two special cases will be examined separately.
6. **Student's *t* random variable.** Blattberg and Gonedes (1974) proposed the Student's *t* random variable as a model for the daily return on common stocks. They compared the fit of the stable and the Student's *t* to daily returns on thirty US stocks. The Student's *t* was considered the better fit for two reasons. Firstly, the estimated degrees of freedom increased as the frequency of observation decreased, indicating a trend towards normality. This is contrary to the assumption of the stable-law, where non-normality would persist under addition. Secondly, the log-likelihood of the Student's *t* was always greater than that of the stable, the justification for this argument draws on an interpretation of log-likelihood ratios as log-odds.
7. **Generalised Error random variable.** The generalised error distribution (GED) has been used to capture leptokurtosis in the conditional returns in conjunction with (G)ARCH models by Nelson (1991), and in conjunction with stochastic volatility by Liesenfeld and Jung (2000).

Most of the analyses quoted in this section refer to the unconditional distribution of asset returns: Kon (1984), Venkataraman (1997), Theodossiou (1998), Blattberg and Gonedes (1974), and Lillestøl (2000). Nelson (1991) modelled the conditional distribution of asset returns. Modelling the conditional heteroscedasticity of asset returns captures some of the observed leptokurtosis in the unconditional distribution of returns. In this analysis, both unconditional and conditional density functions will be estimated, but the emphasis will be on the latter.

4. Analysis of Daily Brent Crude 1991 to 2006

4. 1 Exploratory Data Analysis

The data is divided into three approximately five year sections to examine whether there are any structural changes in the fifteen years of daily data available. As a later part of the analysis involves out of sample forecasting, the final year is excluded from the three sections and used only for validation. Some stylised facts describing the returns on oil prices during each section and the whole period are given in Table 1.

Table 1: Summary Statistics on Brent Oil Price Returns

From		1991/04/08	1995/12/08	2000/08/09	1991/04/08
To		1995/12/07	2000/08/08	2005/04/11	2005/04/11
Number		1219	1218	1219	3656
Maximum		0.0653	0.1516	0.1153	0.1516
Minimum		-0.0756	-0.1260	-0.1697	-0.1697
Mean		0.0000	0.0004	0.0005	0.0003
Standard Deviation		0.0151	0.0244	0.0242	0.0217
Skewness		-0.13	0.04	-0.26	-0.10
Excess Kurtosis		1.85	2.64	3.14	3.57
	Lag				
ACF (Returns)	1	0.06*	0.023	0.01	0.023
	2	-0.052*	-0.024	0.03	-0.005
	3	-0.078*	-0.045	0.038	-0.017
	4	-0.018*	0.013	-0.012	-0.001
	5	0.014*	0.024	-0.015	0.006
ACF (Returns^2)	1	-0.003	0.085*	0.052	0.086*
	2	0.132*	0.033*	0.073*	0.081*
	3	0.049*	0.031*	0.072*	0.075*
	4	0.084*	0.013*	0.069*	0.067*
	5	0.133*	0.121*	0.069*	0.118*

Notes: * indicates p-value of Box Ljung statistic is less than 5%.

Summary statistics give a guide to the presence or absence of properties in the data; in marginal cases the guidance may be wrong. Bearing this comment in mind, one can see from the data that the mean return increases in each five year period, with an average of 0.03% per day over the whole period. Volatility measured by standard deviation was higher in the second and third five year periods than in the first. Skewness is small and fluctuates in sign. There is excess kurtosis in each period and this increases in each five year period. In terms of time series structure, the returns show evidence of significant autocorrelations for January 1991 to December 1995. In terms of volatility clustering, the autocorrelations of squared returns are significant for all periods; this is evidence of a GARCH process.

4.2 Data Analysis and Model Comparison

The continuous time literature on oil pricing has suggested two models, namely, the geometric Brownian motion and a mean reversion model. These models are represented by (10) and (11), respectively. The discrete time considers AR(1) with GARCH(1,1) shown in (7) and (8). Note that here we do not consider a GARCH in mean effect. In addition, Section 3.4 introduces six non-Gaussian density functions for the error term. The number of different models resulting from the combinations of processes for the expected return, unconditional or conditional variance, and choice of density function is large. The models were fitted initially to five years data and used to forecast the next year out of sample. The estimation region was increased by a year and the exercise repeated until the data were exhausted. In order to measure the models' forecasting accuracy, the model was used to simulate the next year's returns. These simulations provide an empirical density function of the return conditional on the correct model identification.

4.3 Tests for Interval Forecasting

Two tests are used to test the null hypothesis that the chosen model is a feasible representation of the data. The first is due to Christoffersen (1998); this is a test for interval forecasts. The test combines both coverage (the proportion within the interval) and independence between occurrences of membership of the interval. The likelihood ratio statistic is

$$LR_{Chris} = -2\left(\ln\left((1-p)^{n_0}\ p^{n_1}\right) - \ln\left((1-\hat{\pi})^{n_0}\ \hat{\pi}^{n_1}\right)\right) \sim \chi_2^2, \qquad (13)$$

where p is the theoretical probability of falling within the interval (e.g. 50%), n_0 and n_1 are the number of times the observation falls outside the interval and within it, and $\hat{\pi} = \frac{n_1}{n_0 + n_1}$.

The second test is due to Berkowitz (2001) and is concerned with the forecast density. If the forecast density of returns is $\hat{f}(u)$ and the cumulative density of an observed return, z_t, is $x_t = \int_{-\infty}^{z_t} \hat{f}(u)\,du$, following a suggestion by Diebold, Gunther and Tay (1998), Berkowitz (2001) argues that it is easier to test for normality rather than for an arbitrary density function. Thus, under the null hypothesis that the forecast density has been correctly identified, it follows that $\Phi^{-1}(x_t) \sim N(0,1)$. Again, the likelihood ratio includes coverage and independence. For an autoregressive process of order one, $Y_t = \mu + \rho Y_{t-1} + \varepsilon_t$, the likelihood ratio is

$$LR(\mu,\sigma,\rho) = -\frac{1}{2}\ln(2\pi) - \frac{1}{2}\ln\left(\frac{\sigma^2}{1-\rho^2}\right) - \frac{\left(Y_1 - \mu/(1-\rho)\right)^2}{2\sigma^2/1-\rho^2}$$

$$-\frac{T-1}{2}\ln(2\pi) - \frac{1}{2}\ln(\sigma^2) - \sum_{t=2}^{T}\left(\frac{(Y_t - \mu - \rho Y_{t-1})^2}{2\sigma^2}\right).$$

The Berkowitz statistic is

$$-2\left(LR(0,1,0) - LR(\hat{\mu},\hat{\sigma},\hat{\rho})\right) \sim \chi_3^2 . \tag{14}$$

4.4 Comparison of Models

Each density model combination was fitted to the data over an increasing length of time, starting at five years, then six years for a total of ten origins. At each origin, the estimated parameter values were used to simulate 1000 iterations of 1 to 260 days ahead. There are two sources of randomness in this simulation. The estimation covariance matrix is used to simulate estimation error, thus 1000 draws are made of possible sets of coefficient values, for each of these draws a set of 260 errors are drawn with zero mean and unit variance and the appropriate density function. These simulations provide both the prediction intervals and the predicted density of returns (and prices although all decisions are based on returns).

Table 2 summarises the results from this exercise. The quality of each 1 to 260 day forecast is measured by membership of the prediction interval (where the probabilities of falling within the interval are 0.99, 0.95, 0.90, 0.80, 0.60, 0.50), and by the Berkowitz statistic (14) for the accuracy of the density forecast. To simplify matters, the number of times the null hypothesis (that the data could be generated by the model) was not rejected was counted. It is recognised that the accuracy of the various interval forecasts is not independent; the number of times the interval test is 'passed' is used as a crude scoring measure.

Table 2: The number of times the model/density forecast was not significant at 5% over ten forecast origins, each with a horizon of one year

Density	Model	Interval						Berkowitz	Total
		0.99	0.95	0.9	0.8	0.6	0.5		
Gaussian	Constant variance	5	2	3	4	6	6	2	28
	Mean Reversion (MR)	6	3	4	5	6	7	2	33
	GARCH	6	5	3	3	7	8	4	36
	GARCH + MR	5	4	4	5	7	8	3	36
	GARCH + AR	5	5	5	5	5	8	2	35
2 Gaussians	GARCH	5	6	5	6	6	6	8	42
	GARCH + MR	6	5	7	5	7	6	5	41
	GARCH + AR	5	7	4	4	6	6	5	37
Gaussian - Laplace Mixture	GARCH	6	7	7	5	6	6	5	42
	GARCH + MR	4	4	8	6	6	4	6	38
	GARCH + AR	5	5	5	3	6	5	7	36
NIG	GARCH	7	7	9	6	7	7	8	51
	GARCH + MR	7	8	8	4	5	6	7	45
	GARCH + AR	5	8	6	5	8	7	7	46
	GARCH + MR + AR	*7*	*7*	*5*	*7*	*7*	*7*	*8*	*48*
SGST	GARCH	5	5	6	5	6	4	6	37
	GARCH + MR	6	6	8	5	6	6	7	44
	GARCH + AR	6	7	7	7	6	7	7	47
Student's t	GARCH	5	6	5	4	5	4	6	35
	GARCH + MR	5	7	8	7	6	6	8	47
	GARCH + AR	5	7	6	5	6	6	6	41
GED	GARCH	5	7	7	4	4	5	6	38
	GARCH + MR	4	6	3	5	6	5	7	36
	GARCH + AR	5	5	6	6	7	6	7	42

The total score (out of 70) shows that the constant variance (geometric Brownian motion) and the mean reversion model with a Gaussian density do not predict the intervals well or match the empirical density of returns. The introduction of the GARCH effect makes some improvement to the total score. Introduction of mean reversion or an autoregressive term makes little difference to the accuracy score. Changing from Gaussian density to the alternatives discussed has an effect. It is clear that density functions with the ability to catch fatter tails lead to more accurate forecasts. Distinguishing between the contenders is more difficult, comparing the average score of the densities with that of the Gaussian; the more convincing contenders are the mixture of two Gaussians, the Student's *t*, the SGST, and the NIG. The evidence for mean reversion or an autoregressive term is mixed; two of these four score best with GARCH only, one scores better with mean reversion, the other with an autoregressive term.

The evidence for GARCH is strong: all the estimations show significant estimates. For mean reversion, the evidence diminishes as the series gets longer. A selection of estimates is shown in Table 3.

Table 3: Mean reversion estimates evolving over ten years

Forecast	GAUSSIAN Mean Reversion			Gaussian/ GARCH/Mean Reversion			NIG/GARCH/Mean Reversion		
Origin	Long-term mean	reversion rate	p value	Long-term mean	reversion rate	p value	Long-term mean	reversion rate	p value
1996/04/19	17.7	0.0069	0.02	18.0	0.0063	0.05	18.3	0.0054	0.05
1997/04/18	18.0	0.0066	0.01	18.3	0.0066	0.02	17.2	0.0045	0.10
1998/04/17	17.6	0.0066	0.01	18.3	0.0087	0.00	18.6	0.0074	0.00
1999/04/16	16.8	0.0050	0.00	17.7	0.0045	0.04	18.1	0.0044	0.03
2000/04/14	17.8	0.0044	0.01	19.4	0.0033	0.09	17.5	0.0031	0.13
2001/04/13	19.2	0.0032	0.01	19.0	0.0030	0.07	46.0	0.0005	0.75
2002/04/12	19.2	0.0037	0.00	19.7	0.0032	0.04	17.6	0.0022	0.17
2003/04/11	19.9	0.0033	0.01	20.9	0.0022	0.12	21.7	0.0011	0.41
2004/04/09	21.4	0.0026	0.03	22.0	0.0020	0.12	22.3	0.0031	0.01
2005/04/08	27.6	0.0011	0.25	1242.7	0.0001	0.92	13.5	-0.0013	0.21

The mean reversion model is plausible for the earlier part of the data. The long-term mean price hovers around \$18 to \$19. However, once the data trends upwards the reversion rate becomes insignificant (the table is shaded to show insignificant mean reversion). The autoregressive term (not reported) in the most accurate AR model, the SGST, was significant at 10% (8 out of 10 times at 5%) and varied in value between 0.04 and 0.08.

One possibility was to consider mean reversion and an autoregressive term simultaneously. When considered with the NIG density, the total score is 48 (as shown in Table 2). However, there were only three out of ten estimations when both mean reversion and autoregressive terms were both significant at 5%.

As examples of the output from the models, Figures 3 and 4 contrast the prediction intervals for returns for the Gaussian mean reversion model and the NIG GARCH model. The NIG GARCH models capture the tail behaviour far better than the mean reversion model. The same information is presented as prediction intervals for prices in Figures 5-8. The prediction intervals show how the feasible range for the price increases with the forecast horizon. The predicted density of prices one year ahead for each model is shown in Figures 6 and 8. The NIG GARCH model allows for the possibility of higher price rises than the mean reversion model.

Figure 3

Mean reversion model - actual returns and quantiles for

17 May 2001 onwards.

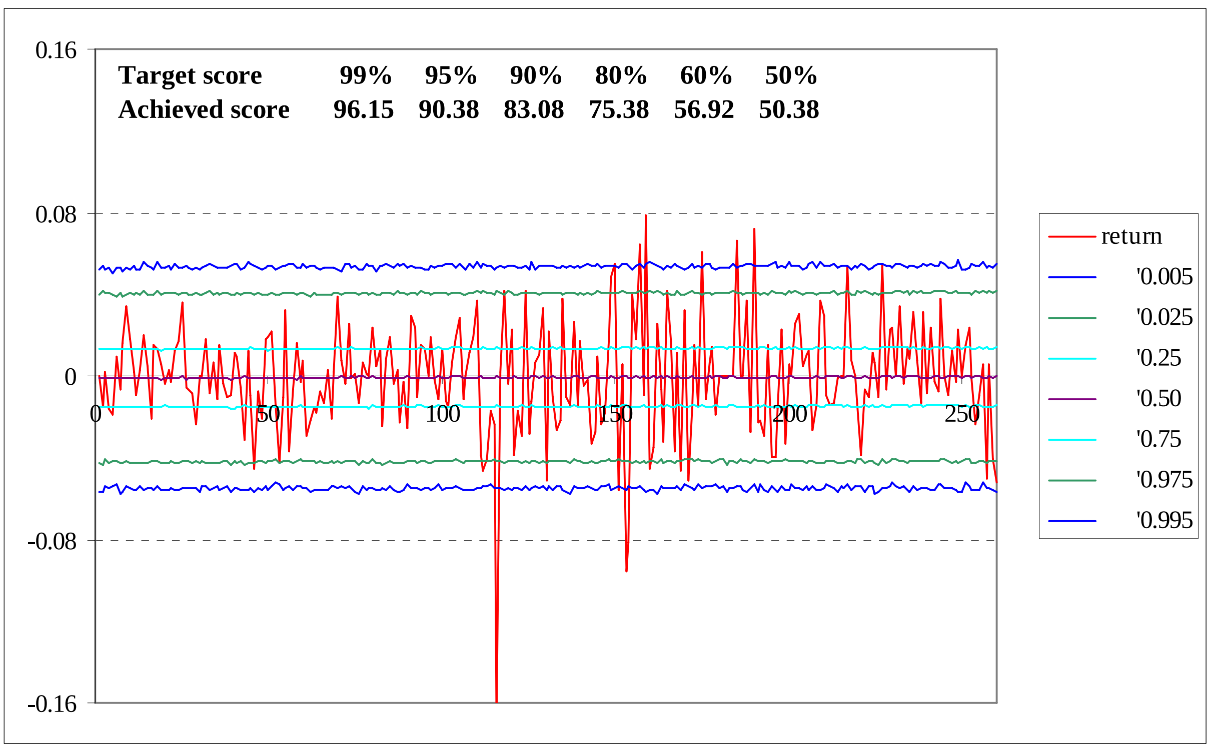

The same procedure was repeated with an estimation region of a fixed five year length. This estimation regime means that the average age of the data is lower, but there are fewer observations for estimation. The total scores (comparable with the right hand column of Table 2) were on average 55% (with a maximum of 77%) of the values in Table 2. Thus, in a tradeoff between more recent data for more up to date estimates of mean returns and more data for better parameter estimation, our research shows that more data is the more effective choice.

Figure 4

NIG GARCH model - actual returns and quantiles for

17 May 2001 onwards.

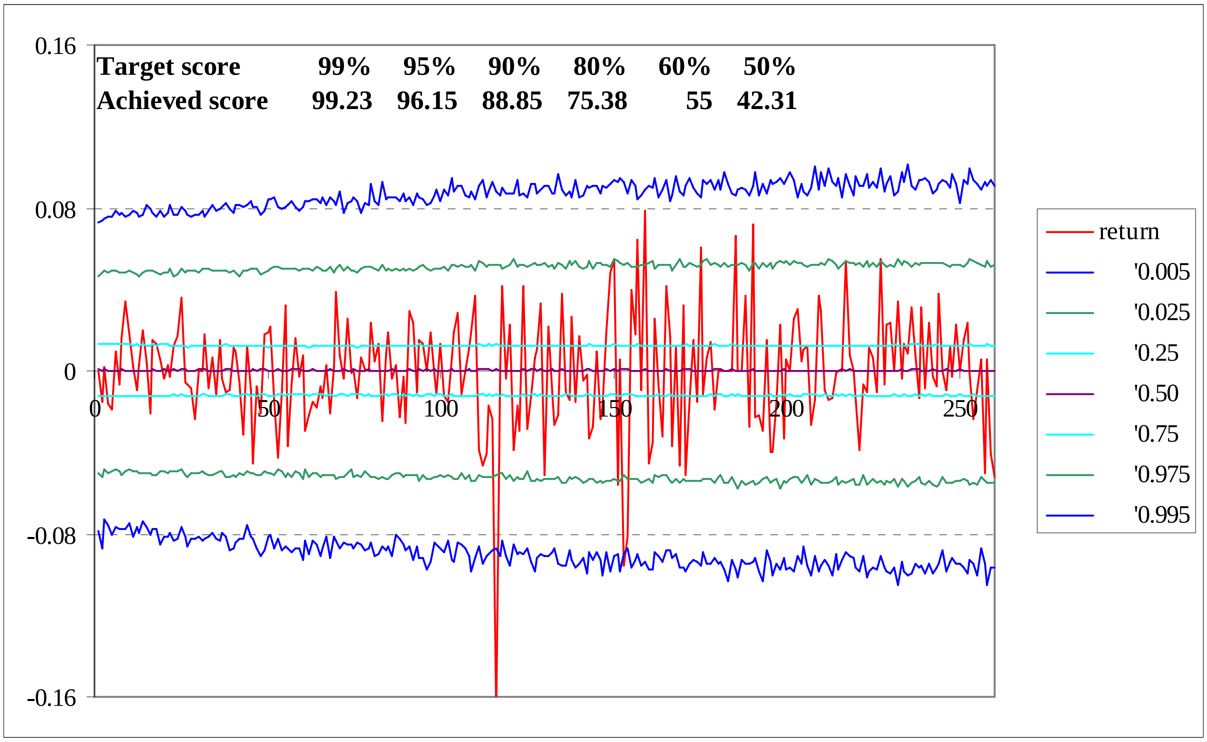

Figure 5

Mean reversion model - actual prices and quantiles for

17 May 2001 onwards.

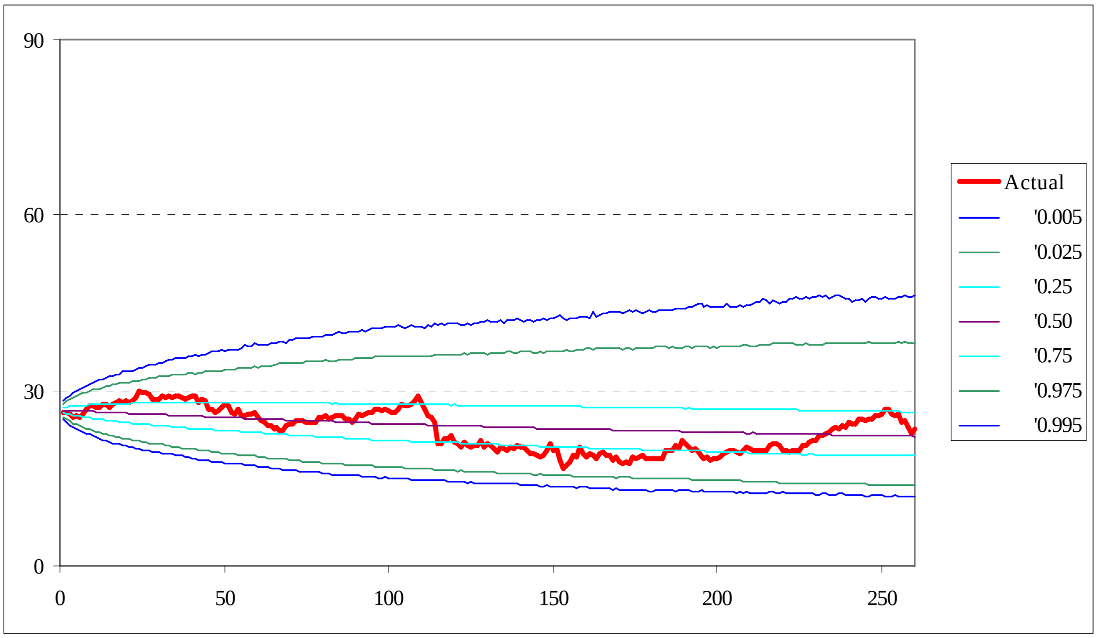

Figure 6

Mean reversion model – predicted density for the price one year ahead.

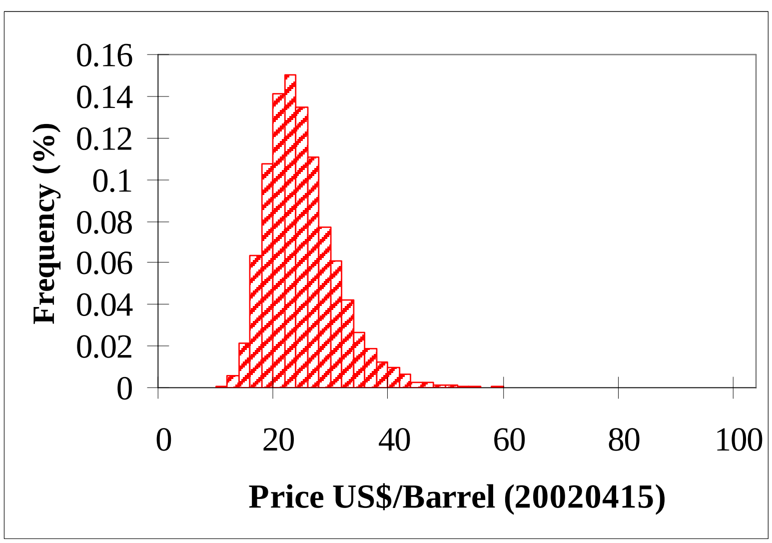

Figure 7

NIG GARCH model - actual prices and quantiles for

17 May 2001 onwards.

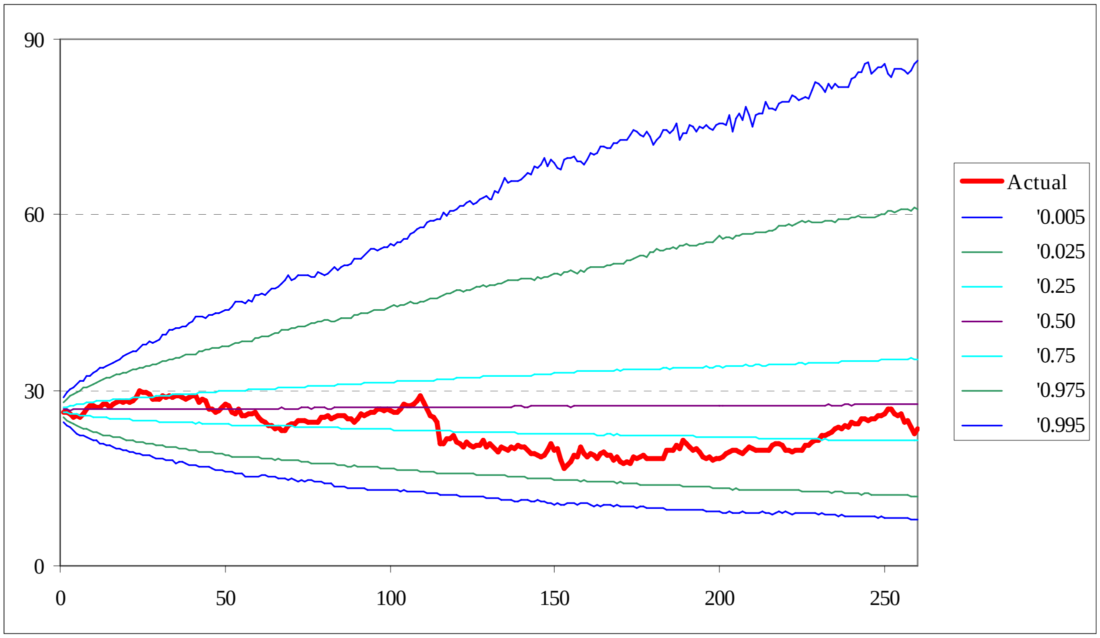

Figure 8

NIG GARCH model - predicted density for the price one year ahead.

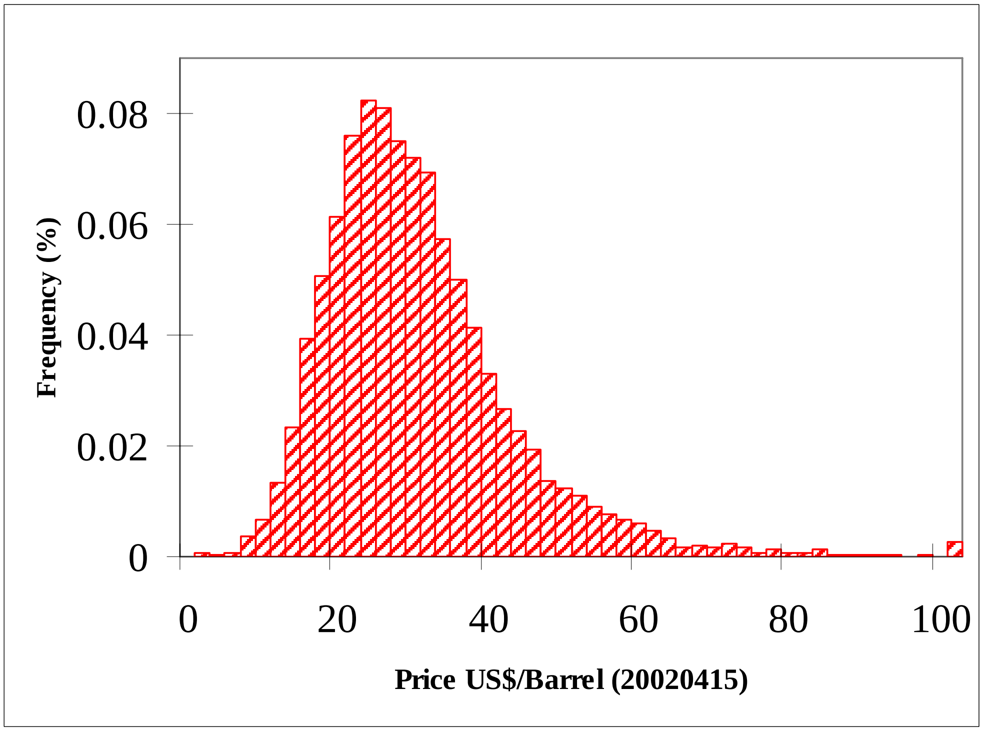

5 Conclusions and Suggestions for Further Research

The main conclusion is that geometric Brownian motion and mean reversion models, on their own, are poor models of oil price returns. Since kurtosis and time varying volatility are important properties of oil price returns, models which include these features provide more accurate density forecasts. The NIG density function appears to be the most accurate of those tested. However, these results are still at a preliminary stage; there is still much work to be done.

Further research into modelling the time series should examine the possibility of introducing time dependent mean returns. An alternative specification would be a time-varying long-term mean price to which oil prices may revert, along the lines of Schwartz and Smith (2000).

Appendix: Density Functions Used in the Analysis

The return random variable *X* has $E(X)=\mu$ and $V(X)=\sigma^2$.

[1] The Gaussian Random Variable:

$$f(x)=\frac{1}{\sigma}\phi\left(\frac{x-\mu}{\sigma}\right), \text{ where } \phi(t)=\frac{e^{-t^2/2}}{\sqrt{2\pi}}.$$

[2] A Mixture of *m* Gaussian Random Variables:

$$f(x)=\sum_{j=1}^{m}\frac{\lambda_j}{\sigma_j}\phi\left(\frac{x-\mu_j}{\sigma_j}\right) \text{ where } \sum_{j=1}^{m}\lambda_j=1\ ;\ \lambda_j>0;\ \sum_{j=1}^{m}\lambda_j\mu_j=\mu \text{ and } \sum_{j=1}^{m}\lambda_j\sigma_j^2=\sigma^2.$$

[3] A Mixture of Gaussian and Laplace Random Variables:

$$f(x)=\frac{\lambda_1}{\sigma_1}\phi\left(\frac{x-\mu_1}{\sigma_1}\right)+\frac{\lambda_2}{2\sigma_2}e^{-\left|\frac{x-\mu_2}{\sigma_2}\right|},$$

where the same conditions apply as in [2] above with *m*=2.

[4] The Normal Inverse Gaussian Random Variable:

$$f(x)=A(\alpha,\beta,\eta,\delta)K_1\left(\delta\alpha q\left(\frac{x-\eta}{\delta}\right)\right)\exp(\beta x)q^{-1}\left(\frac{x-\eta}{\delta}\right),$$

where $q(z)=\sqrt{1+z^2}$, $A(\alpha,\beta,\eta,\delta)=\frac{\alpha\exp(\delta\gamma-\beta\eta)}{\pi}$, $\gamma=\sqrt{\alpha^2-\beta^2}$, and $K_1(\bullet)$ is a modified Bessel function of the third kind of order one.

Given the mean and the variance of the returns, specifying α and β determines the other two parameters. The variance of *X* is given by $V(X)=\sigma^2=\delta\frac{\alpha^2}{\gamma^3}$ which determines δ. The expected value of *X* is $E(X)=\mu=\eta+\delta\frac{\beta}{\gamma}$ which determines η.

[5] The Skewed Generalised Student's *t* Random Variable:

$$f(x)=C\left(1+\tfrac{k}{n-2}\theta^{-k}\left(1+sign(x)\lambda\right)^{-k}\left|\frac{x}{\sigma}\right|^{k}\right)^{\frac{-(n+1)}{k}},$$

where *k, n,* λ, and σ^2 are scaling parameters (*k* > 0, *n* > 2 and $|\lambda|<1$).

The following conditions ensure that *f(x)* is a density function and that the variance is σ^2:

$$C = 0.5kB\left(1/k, n/k\right)^{-\frac{3}{2}} B\left(3/k, (n-2)/k\right)^{\frac{1}{2}} S(\lambda)\sigma^{-1},$$

(where $B(\bullet)$ is the beta function)

$$\theta = \left(\frac{k}{(n-2)}\right)^{\frac{1}{k}} B\left(1/k, n/k\right)^{\frac{1}{2}} B\left(3/k, (n-2)/k\right)^{-\frac{1}{2}} S(\lambda)^{-1}, \text{ and}$$

$$S(\lambda) = \left(1 + 3\lambda^2 - 4\lambda^2 B\left(2/k, (n-1)/k\right)^2 B\left(1/k, n/k\right)^{-1} B\left(3/k, (n-2)/k\right)^{-1}\right)^{\frac{1}{2}}.$$

Relevant special cases are given by further parameter constraints:

Random variable	k	λ	n
Gaussian	2	0	∞
Student's t	2	0	
Generalised error		0	∞
Laplace	1	0	∞

[6] The Student's *t* Random Variable:

$$f(x) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{B\Gamma\left(\frac{\nu}{2}\right)\sqrt{\pi\nu}} \left(1 + \frac{\left(\frac{x}{B}\right)^2}{\nu}\right)^{\frac{-(\nu+1)}{2}},$$

where ν is the degrees of freedom, and the variance is $\left(\frac{\nu}{\nu-2}\right)B^2$, and $\Gamma(\bullet)$ is the gamma function.

[7] The Generalised Error Distribution:

$$f(x) = \frac{k \exp\left(-\frac{1}{2}\left|\frac{x}{C}\right|^k\right)}{C\Gamma\left(1/k\right)2^{\left(1+1/k\right)}}, \quad \text{where} \quad C = \sqrt{2^{-2/k}\frac{\Gamma\left(1/k\right)}{\Gamma\left(3/k\right)}} \quad \textit{and} \quad k > 0.$$

References

[1] Akgiray, V., and G. G. Booth (1988). The Stable Law Model of Stock Returns. Journal of Business & Economic Statistics, 6, 51-57.

[2] Barndorff-Nielsen, O.E. (1997). Normal Inverse Gaussian Process and Stochastic Volatility Modelling. Scandinavian Journal of Statistics, 24, 1-13.

[3] Berkowitz, J. (2001). Testing Density Functions with Applications to Risk Management. Journal of Business & Economic Statistics, 19, 465-474.

[4] Blattberg, R. C., and N. J. Gonedes (1974). A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices. Journal of Business, 47, 244-280.

[5] Brennan, M. J., and E. S. Schwartz (1995). Evaluating Natural Resource Investments. Journal of Business, 58, 135-157.

[6] Cabedo, J. D, and I. Moya (2003). Estimating Oil Price 'Value at Risk' Using the Historical Simulation Approach. Energy Economics, 25, 239-253.

[7] Christoffersen, P. F. (1998). Evaluating Interval Forecasts. International Economic Review, 39, 841-862.

[8] Cortazar G., and L. Naranjo (2006). An N-factor Gaussian Model of Oil Futures Prices. Journal of Futures Markets, 26, 243-68.

[9] Diebold, F. X., Gunther, T. A., and A. S. Tay (1998). Evaluating Density Forecasts with Applications to Financial Risk Management. International Economic Review, 39, 863-883.

[10] Higham, D. J. (2001). An Algorithmic Introduction to Numerical Simulation of Stochastic Differential Equations. SIAM Review, 43, 525-546.

[11] Kon, S. J. (1984). Models of Stock Returns - a Comparison. Journal of Finance, 39, 147-165.

[12] Liesenfeld, R., and R. C. Jung (2000). Stochastic Volatility Models: Conditional Normality Versus Heavy-tailed Distributions. Journal of Applied Econometrics, 15, 137-160.

[13] Lillestøl, J. (2000). Risk Analysis and the NIG Distribution. Journal of Risk, 2, 41-56.

[14] Mandlebrot, B. (1963). The Variation of Certain Speculative Prices. Journal of Business, 36, 394-419.

[15] Mauleon, I., and J. Perote (2000). Testing Densities with Financial Data: an Empirical Comparison of the Edgeworth-Sargan Density to the Student t. European Journal of Finance, 6, 225-239.

[16] Mills, T. C. (1995). Modelling Skewness and Kurtosis in the London Stock Exchange-FTSE Index Return Distributions. Statistician, 44, 323-332.

[17] McDonald, J. B., and W. K. Newey (1988). Partially Adaptive Estimation of Regression Models via the Generalised t Distribution. Econometric Theory, 4, 428-457.

[18] Muralidhar, K. (1993). The Bootstrap Approach for Testing Skewness Persistence. Management Science, 39, 487-491.

[19] Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns. Econometrica, 59, 347-370.

[20] Osborne, M. F. M. (1959). Brownian Motion in the Stock Market. Operations Research, 7, 145-173. Reprinted (1964) in The Random Character of Stock Market Prices, 100-128, P. Cootner (Ed.), MIT Press, Cambridge, Mass.

[21] Panas, E., and V. Nini (2000). Are Oil Markets Chaotic? A Non-linear Dynamic Analysis. Energy Economics, 22, 549-568.

[22] Peiro, A. (1999). Skewness in Financial Returns. Journal of Banking and Finance, 23, 847- 862.

[23] Rydberg, T. H. (1999). Generalised Hyperbolic Diffusion Processes with Applications in Finance. Mathematical Finance, 9, 183-201.

[24] Schwartz, E. S. (1997). The Stochastic Behaviour of Commodity Prices: Implications for Valuation and Hedging. Journal of Finance, 52, 923-973.

[25] Schwartz, E. S., and J. E. Smith (2000). Short-term Variation and Long-term Dynamics in Commodity Prices. Management Science, 46, 893-911.

[26] Singleton, J. C., and J. Wingender (1986). Skewness Persistence in Common Stock Returns. Journal of Financial and Quantitative Analysis, 21, 335-341.

[27] Stevens, P. (2005). Oil Markets. Oxford Review of Economic Policy, 21, 19-42.

[28] Theodossiou, P. (1998). Financial Data and the Skewed Generalised t Distribution. Management Science, 44, 1650-1661.

[29] Venkataraman, S. (1997). Value at Risk for a Mixture of Normal Distributions: The Use of Quasi-Bayesian Estimation Techniques. Federal Reserve Bank of Chicago Economic Perspectives, March/April, 2-13.