



WP/06/299

IMF Working Paper

Recent Dynamics of Crude Oil Prices

Noureddine Krichene

IMF Working Paper

African Department

Recent Dynamics of Crude Oil Prices

Prepared by Nouredine Krichene¹

Authorized for distribution by Benedicte Vibe Christensen

December 2006

Abstract

This Working Paper should not be reported as representing the views of the IMF.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

Crude oil prices have been on a run-up spree in recent years. Their dynamics were characterized by high volatility, high intensity jumps, and strong upward drift, indicating that oil markets were constantly out-of-equilibrium. An explanation of the oil price process in terms of the underlying fundamentals of oil markets and world economy was provided, viewing pressure on oil prices mainly as a result of rigid crude oil supply and an expanding world demand for crude oil. A change in the oil price process parameters would require a change in the underlying fundamentals. Market expectations, extracted from call and put option prices, anticipated no change, in the short term, in the underlying fundamentals. Markets expected oil prices to remain volatile and jumpy, and with higher probabilities, to rise, rather than fall, above the expected mean.

JEL Classification Numbers: C10 C20 C32 Q4.

Keywords: Characteristic function, crude oil, cumulants, drift, Fourier inversion, jump-diffusion, kurtosis, option pricing, skewness, variance-gamma distribution, volatility.

Author's E-Mail Address: nkrichene@imf.org; nkrichene@hotmail.com; nkrichene@isdb.org

¹ The author expresses deep gratitude to Abbas Mirakhor, Abdoulaye Bio-Tchané, Genevieve Labeyrie, Robert Flood, Saeed Mahyoub, U. Jacoby, A. Kovanen, C. Steinberg (all IMF), and Ibrahim Gharbi (the Islamic Development Bank) for their help with this paper.

Contents

I. Introduction	3
II. Empirical Aspects of futures oil prices during 2002–06.....	4
A. Recent Trends and Descriptive Statistics.....	4
B. Oil Price Time-Varying Volatility	7
III. Crude oil price as a Merton Jump-Diffusion Process	9
A. The Stochastic Differential Equation for the Jump-Diffusion Model.....	9
B. Alternative Methods for Estimating the Jump-Diffusion Model: Maximum Likelihood, Method of Cumulants, and Method of Characteristic Function.....	11
IV. Crude Oil Price as a Variance-Gamma Levy Process	17
V. Option Pricing Using Characteristic Functions	19
VI. Density Forecast of Crude Oil Prices: The Inverse Problem.....	21
VII. Conclusions	23
Tables	
1. Time-Series Properties of the Oil Price	6
2. Jump-Diffusion Model: Parameter Estimates.....	15
3. Parameter Estimates of the VG process 1/.....	19
Figures	
1a. Observed Daily Oil Futures Prices.....	5
1b. Actual and Fitted Oil Price	5
2. Daily Crude Oil Price Returns	7
3a. Implied Volatility.....	8
3b. Estimated GARCH.....	8
Annex	
1. Method of Cumulants of Probability Distributions	24
References	
References.....	26

I. INTRODUCTION

Under the combination of an expanding world demand for crude oil and a tight world crude supply, crude oil prices have been on a run-up spree in recent years. By breaking a record level of US\$78.30 per barrel (bl) on August 7, 2006² and remaining comfortably in the neighborhood of US\$75/bl during part of 2006, crude oil prices had risen to unexpected territory and seemed boundless. While developments in crude oil prices were being followed closely by economic agents, including traders, investors, speculators, and policy makers, not much was known about the stochastic processes driving these prices. Contrary to stock market indices, for which an abundant and advanced modeling literature now exists, crude oil prices, in spite of their importance, have not been subject to extensive modeling research. Knowledge of their underlying stochastic process is highly relevant not only for pricing derivatives and hedging, but also for policy making and short-term forecasting.

This paper addressed the dynamics of daily oil prices during January 2, 2002–July 7, 2006.³ Many striking facts regarding oil markets can be stated at the outset. Foremost, world oil demand pressure kept increasing during this period, causing oil prices to rise by more than threefold, from US\$21.13/bl on January, 2002 to US\$73.76/bl on July, 7, 2006.⁴ Second, the noted ascent in oil prices was not monotonic or smooth; oil prices rose, often to a new record, retreated temporarily, then resumed their move to higher record; their movements were dominated by high intensity jumps, indicating that oil markets were constantly out-of-equilibrium. Third, oil price volatilities were excessively high. Measured by the implied volatility, volatility was in the range of 30 percent, implying that oil markets were facing big uncertainties regarding future price developments and were sensitive to small shocks and to news. Finally, market expectations, extracted from crude oil call and put option prices, were right-skewed. More specifically, markets held higher probabilities for further price increases than price decreases. Moreover, markets seemed to expect large upward jumps in oil prices, as reflected by the price and volume of options at strikes in the range of US\$75–US\$85/bl.

When modeled as a jump-diffusion (J-D) process, oil price dynamics were dominated by the discontinuous Poisson jump component compared to the continuous Gaussian diffusion component, showing that oil markets were constantly out-of-equilibrium during the sample period and were sensitive to demand and supply shocks and to news. While the variance of the diffusion component was high and significant, it was surpassed by a still higher and significant variance of the jump component. Both variances, together, illustrated the high

² When BP shut down the Prudhoe Bay field in Alaska for pipeline maintenance.

³ Futures' contracts on Brent, three-month delivery; the sample contains 1130 observations. The source is Reuters.

⁴ The recessionary effect of high oil prices has been studied by Hamilton (1983). A considerable literature thereafter has dealt with the relationship between oil shocks and real GDP. By causing a general increase in the price level, an oil shock, *ceteris paribus*, reduces real cash balances and therefore aggregate demand.

volatility of the oil markets. The drift of the diffusion component was, however, very high and significant, indicating that oil prices were strongly influenced by an upward trend. The mean of the jump component was negative; more specifically, sharp upward jumps in oil prices had a temporary restraining effect on oil demand and were followed by a short-lived sequence of price retreats. The mean of the jump component was, however, outweighed by the drift of the diffusion component which kept prices on a rising trajectory.

Oil prices were also modeled as a Levy process (LP) with a variance-gamma (VG) distribution. The findings were similar to the J-D model. The drift component was positive and highly significant, establishing that oil prices were constantly pulled by an upward trend. The variance of the VG distribution was significant and high. The parameter controlling for the jump process was high and significant, indicating that oil prices were largely dominated by the jump component and oil markets were constantly out-of-equilibrium. The skewness of the VG distribution was negative, indicating that large upward moves in oil prices triggered a temporary depressing effect on world oil demand, translated into a temporary sequence of small negative jumps in oil prices. However, the upward momentum outweighed the small negative jumps. Turning to market expectations, the implied risk-neutral distribution from call and put option prices, assuming a VG process, showed that market participants held higher probabilities for oil prices to rise, rather than fall, above the futures price, and expected oil prices to remain volatile and dominated by a jump process.

The paper is structured as follows. Section I is an introduction. Section II describes the time series properties of oil prices and the empirical distribution of oil price returns. Section III models oil prices as a Merton (1976) J-D process. In Section IV, oil price returns are modeled as a Levy process of the variance-gamma type (Madan and Milne, 1991; and Madan et al., 1998). Section V discusses option pricing in the Fourier space (Carr and Madan, 1999). Section VI presents oil price density forecast based on option prices. Section VII concludes.

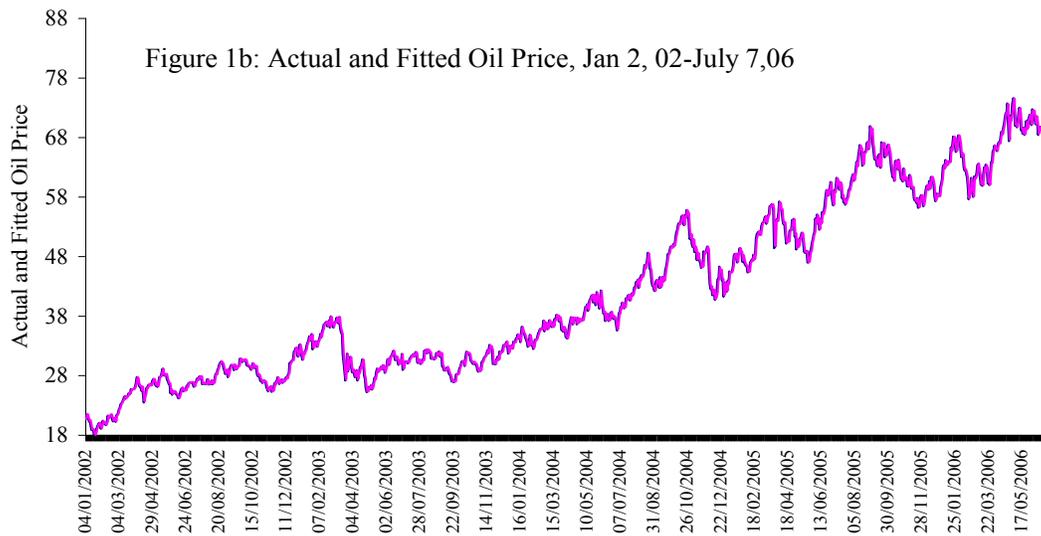
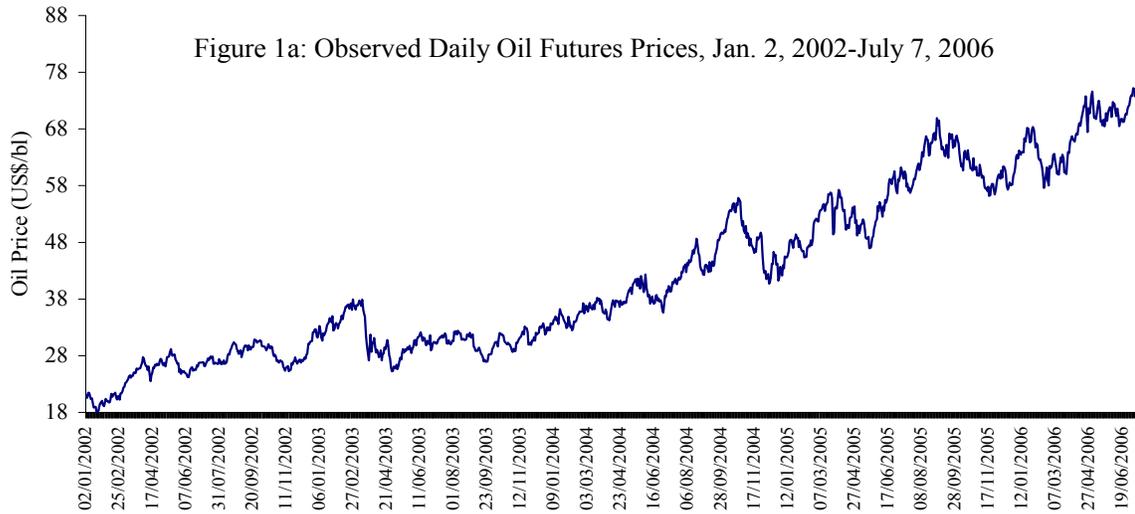
II. EMPIRICAL ASPECTS OF FUTURES OIL PRICES DURING 2002–06

A. Recent Trends and Descriptive Statistics

With a view to concentrating on recent oil prices dynamics, the sample period was chosen to be January 2, 2002–July 7, 2006, containing 1130 daily observations. Figure 1a illustrates the daily behavior of oil prices. It clearly shows that oil prices were moving upward, and have become forecastable. After each peak, oil prices seemed to retreat temporarily then re-trended toward a higher peak. Let S_t be the oil futures price in US\$/bl. An augmented Dickey-Fuller test (Table 1) indicated that S_t possessed a unit root; it was pulled by an upward trend, showing no sign for mean reversion. Changes in S_t , defined as $\Delta S_t = S_t - S_{t-1}$, were, however, stationary. Based on the unit-root test, the dynamics of the oil process were represented by a simple auto-regression of order two (AR2) which yielded good fit and

highly significant coefficients. The graph of this auto-regression (Figure 1b) shows that the AR2 process closely tracks daily oil prices during the sample period.

Figure 1. Daily Oil Futures Prices, January 2, 2002-July 7, 2006.



Source: Reuters.

Table 1. Time-Series Properties of the Oil Price

Augmented Dickey-Fuller Unit-Root Test on Oil Price.
Null Hypothesis: S_t has a unit root. Augmented Dickey-Fuller test statistic=-0.52; probability value=0.88. Test critical values: 1% level (-3.44); 5% level (-2.86); 10% level (-2.57).
Null Hypothesis: $\Delta S_t = S_t - S_{t-1}$ has a unit root. Augmented Dickey-Fuller test statistic=-35.98; probability value=0.00. Test critical values: 1% level (-3.44); 5% level (-2.86); 10% level (-2.57).
Oil Price as a Second-Order Auto-Regressive (AR2) Process $S_t = 0.93S_{t-1} + 0.07S_{t-2} + 0.09, R^2 = 0.99, DW = 2.05.$ <small style="margin-left: 100px;">$t=31.3$ $t=2.3$ $t=0.98$</small>

These price dynamics translated the underlying fundamentals of the oil markets.⁵ In spite of rising prices, world crude oil supply was rigid at about 84 million barrels per day (mbd) for most of the sample period.⁶ World crude oil demand was, however, strongly stimulated by a world economy growing at 4–5 percent per year during 2002–06, excessively low nominal interest rates, and sharply depreciated U.S. dollar.^{7 8} In addition, world crude oil demand is known to be highly price inelastic in the short term. Demand price elasticity ranges between -0.01 and -0.04 (See IMF WP/06/62). More specifically, significant increase in oil price would have only a small negative impact on oil demand. World crude oil demand is also known to have high income elasticity. If the technical coefficient between crude oil and real GDP is fixed in the short term, then income elasticity could be close to one. Econometric estimates, however, show that short-run income elasticity ranges between 0.2 and 0.4. The rigidity of crude supply, combined with an expanding world demand for crude oil, has resulted in growing demand-supply imbalances. Given the price inelasticities of both oil demand and supply, any small excess demand (supply) for oil would require large changes in oil prices to clear oil markets.

⁵ Investors and speculators, through opening and closing positions on the futures markets, affect price dynamics and increase price volatility. However, their role is limited to the short run. Given the sample period under study, underlying fundamentals were key determinants of the oil price process. Incidentally, the IMF World Economic Outlook, September 2006, could not establish evidence for a long-term effect of speculation on oil prices.

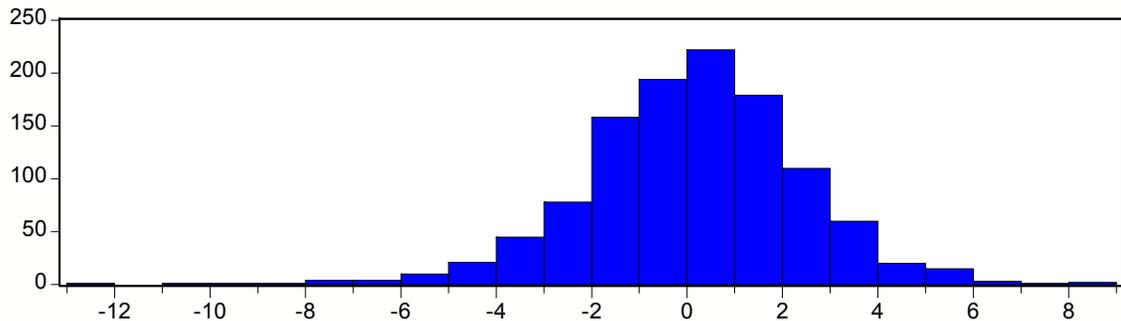
⁶ See, for instance, The International Energy Agency, Oil Market Report, September, 2006. This figure includes natural gas liquids.

⁷ World economy was reported to have grown at about 4-5 percent in real terms during 2002-2006. See International Monetary Fund, World Economic Outlook, September, 2006.

⁸ The relationship between oil prices and economic fundamentals was studied in an IMF working paper (WP/06/62). Besides estimating demand and supply functions for crude oil, the paper analyzed the influence of monetary policy on crude oil prices.

More about oil price dynamics is uncovered by analyzing the log-price return defined as $x_t = \Delta \log S_t = \log S_t - \log S_{t-1}$. The graph for these changes (Figure 2) shows that large jumps in crude oil prices were frequent and had a relatively high probability. Although the mode was around 1–2 percent, daily changes in the range of 5–7 percent were not uncommon.⁹ The empirical distribution had a large dispersion, with standard deviation estimated at 2.29 (annualized to 36.3 percent). The distribution was left-skewed, implying that downward jumps of smaller size were more frequent than upward jumps of larger size; as the mean was positive and high, smaller jumps were outweighed by larger jumps. The distribution had also fat tails, meaning that large jumps tended to occur more frequently than in the normal case. These empirical findings about daily oil futures prices were typical of financial time series as noted in Clark (1973), Fama (1965), and Mandelbrot (1963). These facts suggested modeling the oil price process as a jump-diffusion or, in a more general way, as a Levy process (Cont and Tankov, 2004).

Figure 2. Daily Crude Oil Price Returns Distribution, Jan 2, 02-July 7, 06.



Descriptive statistics: mean = 0.116; standard deviation=2.29; skewness=-0.39; kurtosis=4.79; Jarque-Bera normality statistics=179.3, probability-value=0.0.

B. Oil Price Time-Varying Volatility

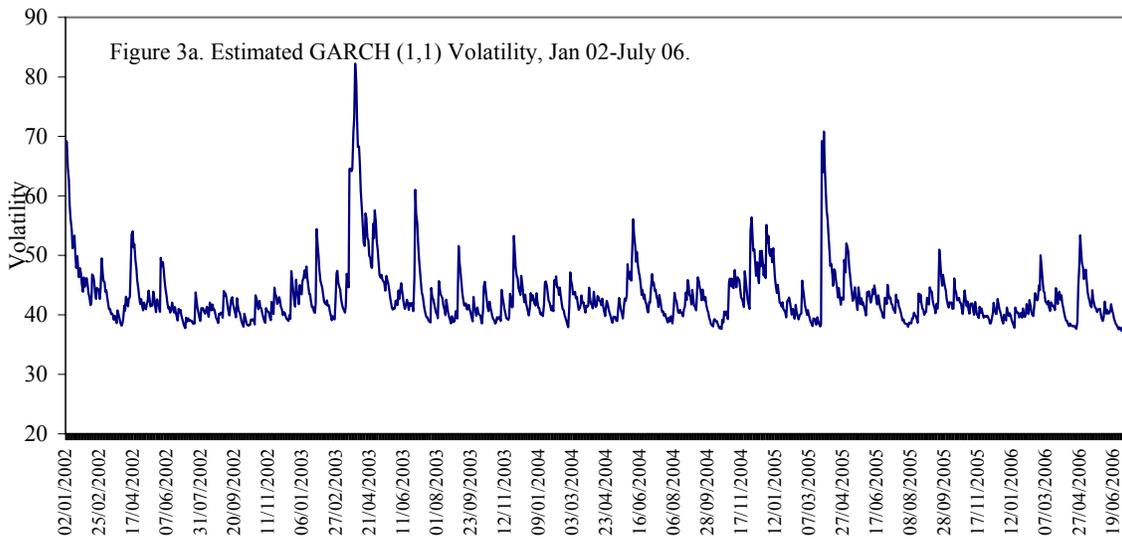
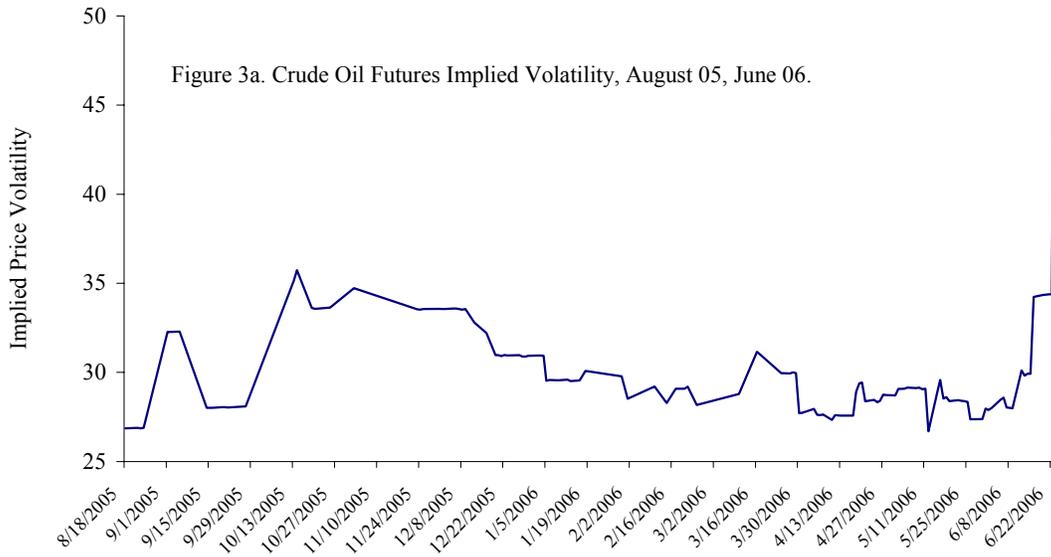
Price volatility is an important parameter in financial markets to which market participants pay a great deal of attention. Volatility measures uncertainty and also sensitivity of prices to news and shocks, and is a key parameter in option pricing. Two types of volatilities are studied here: the implied volatility from crude oil call options,¹⁰ and volatility computed by a GARCH (1, 1) model. Both measures point to high volatility in futures prices.

Data on implied volatility for Brent futures options for August 2005–June 2006 indicated that oil price volatility was high (Figure 3a). While averaging 30 percent, volatility often surged

⁹ The frequency of jumps exceeding ± 3 percent was estimated from the sample at 23 percent.

¹⁰ Implied volatility is the volatility which equates the Black-Scholes (1973) call option pricing formula with the call option's market value.

Figure 3. Implied and GARCH (1,1) Volatility.



Source: Bloomberg for Implied Volatility.

to 34–35 percent, indicating that oil markets were experiencing big uncertainty regarding expected price developments, and were highly sensitive to small shocks and news. Volatility pattern shows volatility clustering during rising pressure on oil prices and volatility decline during reduced pressure on oil prices. High volatility increases speculative demand for futures contracts, which in turn leads to higher volatility and volatility clustering.

Volatility was also computed using a GARCH model for data on daily oil futures prices covering January 2, 2002–July 7, 2006 (Figure 3b). The oil price return was defined as: $x_t = \Delta \log S_t = \log S_t - \log S_{t-1}$.¹¹ The fitting of the GARCH model showed high price volatility, periods of volatility clustering, followed by some reversion to a mean volatility estimated at 43 percent. GARCH volatility was rising during periods of large price shocks, increasing speculation and leading to volatility clustering; it was, however, receding during periods of price retreat. It corroborated the observed implied volatility, namely oil markets were constantly experiencing large uncertainties and were impacted by frequent shocks.

III. CRUDE OIL PRICE AS A MERTON JUMP-DIFFUSION PROCESS

A. The Stochastic Differential Equation for the Jump-Diffusion Model

Based on the empirical findings of the previous section, namely the presence of skewness and kurtosis in the empirical distribution of oil price returns, an adequate model for oil prices would be a jump-diffusion model. In fact, Merton (1976), recognizing the presence of jumps in asset prices and for more accurate option pricing, proposed modeling these prices as a jump-diffusion process instead of a pure diffusion model. Moreover, it is well-known that short-term options have market implied volatilities that exhibit a significant skew across strikes. In this connection, Bakshi et al. (1997) argued that pure diffusion based models have difficulties explaining the smile effect in short-dated option prices and emphasized the importance of adding a jump component in modeling asset price dynamics. In the same vein, Bates (1996) noted that diffusion based stochastic volatility models cannot explain skewness in implied volatilities, except under implausible values for the model's parameters. Models with jumps generically lead to significant skews for short-term maturities. More generally, adding jumps to returns in diffusion based stochastic volatility model, the so-obtained model can generate sufficient variability and asymmetry in the short-term returns to match implied volatility skews for short-term maturities.

¹¹ GARCH stands for Generalized Autoregressive Conditional Heteroskedasticity. A GARCH (1,1) model is defined as follows: The mean equation: $x_t = \Delta \log S_t = c + \varepsilon_t$, $\varepsilon_t \sim N(0, \sigma_t^2)$

The conditional variance equation: $\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$, where $\sigma_t^2 = E(\varepsilon_t^2)$

Accordingly, the continuous-time stochastic process driving crude oil prices can be stated as a J-D process given by a stochastic differential equation

$$\frac{dS_t}{S_t} = \alpha dt + \sigma dB_t + (\exp(J_t) - 1) dN_t$$

S_t denotes the crude oil price, α is the instantaneous return, and σ^2 is the instantaneous variance. The continuous component is given by a standard Brownian motion, B_t , distributed as $dB_t \sim N(0, dt)$. The discontinuities of the price process are described by a Poisson counter N_t , characterized by its intensity, λ , and jump size, J_t . The Brownian motion and the Poisson process are independent. The intensity of the Poisson process describes the mean number of arrivals of abnormal information per unit of time and is expressed as $\text{Prob}[\Delta N_t = 1] = \lambda dt$, and $\text{Prob}[\Delta N_t = 0] = 1 - \lambda dt$. When abnormal information arrives, crude oil price jumps from S_{t-} to $S_t = \exp(J_t) S_{t-}$. The percentage change is measured by $(\exp(J_t) - 1)$. The jump size, J_t , is independent of B_t and N_t , and is assumed to be normally distributed: $J_t \sim N(\beta, \delta^2)$. Letting $X_t = \log(S_t)$ and using Ito's lemma, the log price return process becomes

$$dX_t = \left(\alpha - \frac{1}{2} \sigma^2 \right) dt + \sigma dB_t + J_t dN_t = \mu dt + \sigma dB_t + J_t dN_t$$

where $\mu = \left(\alpha - \frac{1}{2} \sigma^2 \right)$.¹² The parameter vector associated with the price process is therefore $\theta = (\mu, \sigma^2, \lambda, \beta, \delta^2)$. Discretized over $(t, t + \Delta)$, the model takes the form:

$$\Delta X_t = \mu \Delta + \sigma \Delta B_t + \sum_{i=0}^{\Delta N_t} J_i$$

Where $\Delta B_t = B_{t+\Delta} - B_t \sim N(0, \Delta)$, and $\Delta N_t = N_{t+\Delta} - N_t$ is the actual number of jumps occurring during the time interval $(t, t + \Delta)$, and J_i are independently and identically distributed as $J_i \sim N(\beta, \delta^2)$. The log-return, $x_t = \Delta X_t$, includes therefore the sum of two independent components: a diffusion component with drift and a jump component. Its probability density is a convolution of two independent random variables and can be expressed as¹³

¹² A solution to this SDE can be written as $S_T = S(0) \exp \left[\left(\alpha - \frac{1}{2} \sigma^2 \right) T + \sigma B_T + \sum_{i=0}^{N_T} J_i \right]$

¹³ Ball and Torous (1985) modeled the jump component in Merton's model as a Bernoulli process. In this respect, either one or no abnormal event occurs during the time interval $(t, t + \Delta)$, with $\text{Prob}[\text{one abnormal}$

(continued...)

$$f(x) = \sum_{i=0}^{\infty} \frac{(\lambda\Delta)^n e^{-\lambda\Delta}}{n!} \left[\frac{1}{\sqrt{2\pi(\sigma^2\Delta + n\delta^2)}} \exp\left(-\frac{(x - \mu\Delta - n\beta)^2}{2(\sigma^2\Delta + n\delta^2)}\right) \right]$$

With $n = 0, 1, 2, \dots$. Putting $\Delta = 1$, i.e., the time interval is $(t, t + 1)$, the density function becomes

$$f(x) = \sum_{i=0}^{\infty} \frac{(\lambda)^n e^{-\lambda}}{n!} \left[\frac{1}{\sqrt{2\pi(\sigma^2 + n\delta^2)}} \exp\left(-\frac{(x - \mu - n\beta)^2}{2(\sigma^2 + n\delta^2)}\right) \right]$$

B. Alternative Methods for Estimating the Jump-Diffusion Model: Maximum Likelihood, Method of Cumulants, and Method of Characteristic Function

1. The maximum likelihood method: Let $x = \{x_1, x_2, \dots, x_T\}$ be an observed sample of log returns, the log-likelihood function can be expressed as:

$$L(\theta; x) = -T\lambda - \frac{T}{2} \ln(2\pi) + \sum_{t=1}^T \ln \left[\sum_{j=0}^{\infty} \frac{\lambda^j}{j!} \frac{1}{\sqrt{\sigma^2 + n\delta^2}} \exp\left(-\frac{(x_t - \mu - n\beta)^2}{2(\sigma^2 + n\delta^2)}\right) \right]$$

Application of the maximum likelihood (ML) method for estimating the J-D model has met with difficulties arising mainly from the identification of the jump parameter and instability of parameter estimates. Nonetheless, Ball and Torous (1983) applied directly the ML method by truncating the number of jumps at $n = 15$. Ball and Torous (1985) and Jorion (1988) applied the ML method by assuming a Bernoulli process for the jump component. While the ML estimates achieve the lower bound for Cramer-Rao efficiency criterion, difficulties with the likelihood function arising from computational tractability, un-boundedness over the parameter space, and instability of parameters, have led researchers to explore alternative estimation methods, based essentially on the method of moments.

2. The method of cumulants (See Annex I): Press (1967) used the method of cumulants as described in Kendall and Stuart (1977) to estimate the J-D model. Define the characteristic function (CF) of X_t as $\phi_x(u) = E[\exp(iuX_t)] = \int \exp(iuX_t) f(X_t) dX_t$, where $f(X_t)$ is the probability density function of X_t , u is the transform variable,

event] = $\lambda\Delta$, Prob[no abnormal event] = $1 - \lambda\Delta$, and Prob[more than one abnormal event] = 0. The density function for the log-return becomes: $f(x) = N(\mu\Delta + n\beta, \sigma^2\Delta + n\delta^2)(\lambda\Delta) + N(\mu\Delta, \sigma^2\Delta)(1 - \lambda\Delta)$

and $\sqrt{-1} = i$.¹⁴ The cumulants of X_t , denoted by κ_n , $n = 0, 1, 2, \dots$, are the coefficients in the power series expansion of the logarithm of the CF of X_t , expressed as:

$$\ln \phi(u) = \sum_{n=1}^{\infty} \kappa_n \frac{(iu)^n}{n!} = 1 + \kappa_1 \frac{(iu)}{1!} + \kappa_2 \frac{(iu)^2}{2!} + \dots + \kappa_n \frac{(iu)^n}{n!} + \dots$$

Noting that the CF for the jump-diffusion process is given by:¹⁵

$$\phi_{X_1}(u) = \exp \left[-\frac{\sigma^2 u^2}{2} + i\mu u + \lambda \left(\exp \left(i\beta u - \frac{\delta^2 u^2}{2} \right) - 1 \right) \right];$$

It follows that the first four cumulants of the J-D process are $\kappa_1 = \mu + \lambda\beta$,

$\kappa_2 = \sigma^2 + \lambda\delta^2 + \lambda\beta^2$, $\kappa_3 = \lambda\beta(3\delta^2 + \beta^2)$, $\kappa_4 = \lambda(3\delta^4 + 6\beta^2\delta^2 + \beta^4)$. Obviously, the cumulants enable to recover the parameters of the J-D process from sample moments. Press (1967), in order to avoid using higher order cumulants, imposed the restriction $\mu = 0$ and

derived the following relations: $\hat{\beta}^4 - 2\frac{\kappa_3}{\kappa_1}\hat{\beta}^2 + \frac{3\kappa_4}{2\kappa_1}\hat{\beta} - \frac{\kappa_3^2}{2\kappa_1^2} = 0$, $\hat{\lambda} = \frac{\kappa_1}{\hat{\beta}}$, $\hat{\delta}^2 = \frac{\kappa_3 - \hat{\beta}^2\kappa_1}{3\kappa_1}$,

$\hat{\sigma}^2 = \kappa_2 - \frac{\kappa_1}{\hat{\beta}} \left(\hat{\beta}^2 + \frac{\kappa_3 - \hat{\beta}^2\kappa_1}{3\kappa_1} \right)$. Press' estimates were often wrong-signed and not plausible.

Beckers (1981) adopted the same method as Press, however, setting β , instead of μ , to zero.

Using sixth order cumulant, his cumulant equations yielded the following system: $\hat{\mu} = \kappa_1$,

$$\hat{\lambda} = \frac{25\kappa_4^3}{3\kappa_6^2}, \quad \hat{\delta}^2 = \frac{\kappa_6}{5\kappa_4}, \quad \text{and} \quad \hat{\sigma}^2 = \kappa_2 - \frac{5\kappa_4^2}{3\kappa_6}$$

Beckers' estimates improved those of Press, yet they were not free of anomalies. Ball and Torous (1983), using a Bernoulli, instead of a Poisson, jump process and maintaining Beckers' restriction, i.e. $\beta = 0$, derived the following cumulant equations: $\kappa_1 = \mu$,

$\kappa_2 = \sigma^2 + \lambda\delta^2$, $\kappa_3 = 0$, $\kappa_4 = 3\delta^2\lambda(1-\lambda)$, $\kappa_5 = 0$, and $\kappa_6 = 15\delta^6\lambda(1-\lambda)(1-2\lambda)$. Again by

equating with population cumulants, they obtained estimators $\hat{\mu}$, $\hat{\lambda}$, $\hat{\sigma}^2$, and $\hat{\delta}^2$ given by:

$$\hat{\mu} = \kappa_1, \quad \hat{\lambda} = \left(1 \pm \sqrt{3\kappa^* / (3\kappa^* + 100)} \right) / 2, \quad \text{where} \quad \kappa^* = (\kappa_6 / \kappa_4)^2, \quad \hat{\sigma}^2 = \kappa_2 - \hat{\lambda}\hat{\delta}^2, \quad \text{and}$$

$$\hat{\delta}^2 = \kappa_6 / \left(\kappa_4 \left(5(1-2\hat{\lambda}) \right) \right).$$

¹⁴ The characteristic function $\phi_X(u)$ is related to the moment generating function $G_X(u)$

$G_X(u) = E \left[\exp(uX_t) \right] = \int \exp(uX_t) dF(X_t)$ by a change of the transform variable $u \rightarrow -iu$, namely $G_X(iu) = \phi_X(u)$, and $G_X(u) = \phi_X(-iu)$.

¹⁵ See, for instance, Madan and Seneta (1987), and Cont and Tankov (2004).

Das and Sundaram (1999) used the method of moments to estimate the J-D model. Denoting the log-price return by x_t , and assuming that the jump size J_t is distributed as

$J \sim N(\beta, \delta^2)$, they showed that the first moments of the J-D process are given by the

following equations which they used to estimate the model's parameters; however, for the Poisson parameter λ , they imposed a given value.

$$Var[x] = E\left[(x - E(x))^2\right] = (\sigma^2 + \lambda\delta^2)$$

$$skewness(x) = \frac{E\left[(x - E(x))^3\right]}{[Var(x)]^{3/2}} = \frac{\lambda(\beta^3 + 3\beta\delta^2)}{(\sigma^2 + \lambda\delta^2 + \lambda\beta^2)^{3/2}}$$

$$kurtosis(x) = \frac{E\left[(x - E(x))^4\right]}{[Var(x)]^2} = 3 + \frac{\lambda(\beta^4 + 6\beta^2\delta^2 + 3\delta^4)}{(\sigma^2 + \lambda\delta^2 + \lambda\beta^2)^2}$$

3. The method of the characteristic function (CF): As there is a one-to-one correspondence between the CF, $\phi_x(u) = E[\exp(iuX_t)] = \int \exp(iuX_t) f(X_t) dX_t$, and the corresponding probability density, $f(X_t)$, the CF conveys the same information as the probability distribution. Often, the transition density function of a stochastic process may not be available in closed form, while the CF is readily available in closed form. Knowledge of the analytic form of the CF allows estimating the parameters of the process by the method of moments or the empirical CF procedure (ECF).¹⁶ The method of moments computes non-

central moments of any order n as $E[(X_t)^n] = \frac{1}{i^n} \frac{d^n}{du^n} \phi(u)|_{u=0}$. It also enables the

application of the empirical characteristic function method (ECF). In both cases, a General Method of Moments (GMM) procedure is implemented, consisting of minimizing a distance norm between the sample and the theoretical population moments, or the sample CF and the theoretical CF. The exact method of moments consists of estimating the parameter vector

which minimizes the distance $\left\| E(X^n) - \frac{1}{i^n} \frac{\partial^n \phi}{\partial u^n} \Big|_{u=0} \right\|$.¹⁷

¹⁶ Parzen (1962), Feuerverger and Mureika (1977), Feuerverger and McDunnough (1981a and 1981b) suggested the use of the CF to deal with the estimation of density functions. Madan and Seneta (1987) proposed a CF-based approach to estimate the J-D model. In the same vein, Bates (1996), Duffie et al. (2000), Chacko and Viceira (2003), and many other authors have proposed the use of CF for estimating affine J-D models.

¹⁷ Note that $X^n = e^{\log(X^n)} = e^{n \log(X)}$. Therefore,

$$E(X^n) = E\left(e^{\log(X^n)}\right) = E\left(e^{n \log(X)}\right) = \int_{-\infty}^{\infty} e^{n \log(X)} f(X) dX = \phi_{\log(X)}(n). \text{ Namely, for the log-}$$

(continued...)

The ECF method can be described as follows. Suppose $x = \{x_1, x_2, \dots, x_T\}$ is an identically independently distributed realization of the same variable X with density $f(x; \theta)$ and a distribution $F_\theta(x)$. The parameter $\theta \in R^l$ is the parameter of interest with true value θ_0 . It is to be estimated from $x = \{x_1, x_2, \dots, x_T\}$. Define the theoretical CF as: $\phi_\theta(u) = \int e^{iux} f(x; \theta) dx$ and its empirical counterpart (ECF) as

$$\phi_n(u) = \int e^{iux} f_n(x) dx = \frac{1}{T} \sum_1^n \exp(iux_j) = \frac{1}{T} \sum_{j=1}^n \cos(ux_j) + i \frac{1}{T} \sum_{j=1}^n \sin(ux_j)$$

The ECF procedure consists of estimating θ according to the criterion

$$\hat{\theta} = \arg \min_{\theta} (\phi_n - \phi_\theta)' W (\phi_n - \phi_\theta)$$

W is a positive semi-definite matrix. Because the minimization of the distance between the ECF (ϕ_n) and CF (ϕ_θ) over a grid of points in the Fourier domain is equivalent to matching a finite number of moments, the ECF method is in essence equivalent to the Generalized Method of Moments (GMM). Feuerverger (1990) proved that, under some regularity conditions, the resulting estimates can be made to have arbitrarily high asymptotic efficiency provided that the sample of observations is sufficiently large and the grid of points is sufficiently fine and extended. Indeed, ECF estimators have the same consistency and asymptotic efficiency as the GMM estimators. Moreover, when the number of orthogonal conditions exceeds the number of parameters to be estimated ($r > l$), the model is over-identified, in that more orthogonal conditions are used than needed to estimate θ . A test of over-identifying restrictions may be used. In this respect, Hansen (1982) suggested a test of whether all of the sample moments are as close to zero as would be expected if the corresponding population moments were truly zero.

4. Empirical results of the estimation

Based on a sample of daily prices for Brent futures prices described in Section II, alternative methods were used for estimating the J-D model (Table 2). First, assuming a Bernoulli jump process, the ML was applied unrestrictedly, and with restriction on the probability λ of a jump occurring on a trading day given by $\lambda = 0.23$. Second, the method of cumulants was applied consecutively with restrictions $\lambda = 0.23$, $\mu = 0$ (Press, 1967), and $\beta = 0$ (Beckers, 1981), respectively. The third method was the ECF applied unrestrictedly and with restriction

return $x_t = \log(S_t / S_{t-1})$, $E(x^n) = E\left(e^{\log(S_t / S_{t-1})^n}\right) = E(e^{nx}) = \int_{-\infty}^{\infty} e^{nx} f(X) dX = \phi_x(n)$. It follows that

the n -th order moment $E(x^n)$ can be computed by replacing the transform variable u by n in the CF of $x_t = \Delta X_t = \log(S_t / S_{t-1})$.

$\lambda = 0.23$. The three methods yielded parameter estimates that were consistent with the empirical features of oil prices discussed in Section II. They showed pointedly that the dynamics of the oil price process were influenced by both diffusion and jump components; however the jump component was the dominant one. Besides having high intensity, the jump component had a much higher variance than the diffusion component. The high variance of the jump component illustrated the presence of jumps of large magnitude and was in conformity with the excess kurtosis in the empirical distribution of oil price returns. The mean of the jump size tended to be negative, in conformity with the negative skewness of the empirical distribution. This was due to the fact that crude oil prices were not monotonic; they leapt forward, then retreated back in smaller movements before taking a new jump. The drift of the diffusion component was high, in conformity with the observed upward trend in crude oil prices; it illustrated the presence of a force that kept pushing oil prices upward and was able to outweigh the negative mean of the jump component.

Table 2. Jump-Diffusion Model: Parameter Estimates

Methods	Drift μ	Variance σ^2	Intensity λ	Mean β	Variance δ^2
Bernoulli process					
Maximum Likelihood	0.23 (t=3.22)	4.46 (t=20.25)	0.59 (t=1.89)	-1.12 (t=-4.27)	4.47 (t=17.12)
Maximum Likelihood 1/	0.27 (t=3.08)	3.34 (t=14.49)	0.23	-0.68 (t=-1.93)	7.98 (t=6.25)
Cumulants 1/ Press (1967) 2/	0.32 0	1.81 6.54	0.23 0.10	-0.85 1.11	2.78 -13.88
Beckers (1981) 3/ ECF 4/	0.12 0.57 (t=7.96)	3.34 0.54 (t=1.11) 3.45	0.22 4.37 (t=3.49)	0 -0.10 (t=-6.25)	8.62 0.96 (t=5.57) 6.97
ECF 1/ 4/	0.27 (t=16.1)	(t=127)	0.23	-0.52 (t=-3.75)	(t=31.35)

1/ Restriction on $\lambda = 0.23$, computed from the data sample as the frequency of a jump in the crude oil price exceeding ± 3 percent. 2/ Restriction on $\mu = 0$. 3/ Restriction on $\beta = 0$. 4/ The grid for u consists of twenty points: 0.1, 0.2, 0.3,, 1.9, 2.0.

Assuming a Bernoulli jump process, the ML estimates were highly significant and stable. The drift of the diffusion component, estimated at $\hat{\mu} = 0.23$, was very high and significant, showing that oil prices were constantly under pressure to move upward. The variances of the diffusion and jump components were high and significant, $\hat{\sigma}^2 = 4.46$ and $\hat{\delta}^2 = 4.47$, respectively. The variance of the jump component became more important than that of the diffusion component when the jump intensity was restricted to $\lambda = 0.23$. The probability of a jump in the unrestricted case, computed at $\hat{\lambda} = 0.59$, was high and borderline significant. The mean of the jump component, estimated at $\hat{\beta} = -1.12$, was negative and consistent with the negative skewness observed in the data. Oil prices tended to make large moves upward, then started to retreat through a sequence of smaller and frequent negative jumps, until they were shocked again, making new jumps forward. Yet, the significance of the drift of the diffusion

process was such that the smaller negative jumps could not outweigh the strong momentum that kept pushing oil prices upward.

The method of cumulants was applied under alternative restrictions. The restriction $\lambda = 0.23$ yielded results that were similar to the ML under the same restriction. The drift of the diffusion component, estimated at $\hat{\mu} = 0.32$, was very high, showing that oil prices were constantly under pressure to move upward. The variances of the diffusion and jump components, were estimated at $\hat{\sigma}^2 = 1.81$ and $\hat{\delta}^2 = 2.78$, respectively, indicating that the jump component tended to dominate the dynamics of the oil price process. The mean of the jump component, estimated at $\hat{\beta} = -0.85$, was negative and consistent with the negative skewness in oil price returns. Application of the Press (1967) method, with the restriction $\mu = 0$, yielded implausible results for the variance of the jump component, namely $\hat{\delta}^2 = -13.88$. Such an anomaly was not unexpected in the case of Press' method, indicating that the restriction $\mu = 0$, could not be borne by the data, and was in sharp contrast with the strong upward trend in oil prices. In contrast, Beckers' method, with the restriction $\beta = 0$, yielded results which were highly plausible. The drift component of the diffusion, estimated at $\hat{\mu} = 0.12$, was smaller than, say, in the ML case, since $\beta = 0$ implied less influence for the drift of the diffusion, compared to the case when β was negative, to maintain an upward trend in oil prices; it was, however, close to the drift of the AR2 (Table 1) and the actual mean of oil price returns (Figure 2). The variances of the diffusion and jump components were high, $\hat{\sigma}^2 = 3.34$ and $\hat{\delta}^2 = 8.62$, respectively. The variance of the jump component, however, dominated that of the diffusion component. Noticeably, the jump intensity, estimated at $\hat{\lambda} = 0.22$, was quite close to the frequency of jumps in oil prices exceeding ± 3 percent, computed from the data set.

The ECF was applied unrestrictedly and with restriction $\lambda = 0.23$. The drift of the diffusion component, estimated at $\hat{\mu} = 0.57$, was very high and significant. The variance of the diffusion component, $\hat{\sigma}^2 = 0.54$, was not significant, and was dominated by the variance of the jump component, $\hat{\delta}^2 = 0.96$, which was significant. The intensity of the jump process, estimated at $\hat{\lambda} = 4.37$, was high and significant, indicating that the oil price process was characterized by frequent jumps. The mean of the jump component, $\hat{\beta} = -0.1$, was negative, significant, and consistent with skewness in oil price returns. The ECF, applied with the restriction $\lambda = 0.23$, yielded results which were similar to those of the ML using the same restriction. The drift $\hat{\mu} = 0.27$ was positive and significant; the variance of the diffusion, $\hat{\sigma}^2 = 3.34$, was significant; however, it was dominated by the variance of the jump component, $\hat{\delta}^2 = 6.97$, indicating that the jump process played a more important role in oil price dynamics in relation to the diffusion process. The mean of the jump component, $\hat{\beta} = -52$, was negative, significant and consistent with skewness observed in the data.

In sum, parameter estimates from the three methods were fully concordant with the data. They established that the oil price process was dominated by the jump process, with large discontinuities occurring at high intensity, meaning that oil markets were permanently out-of-equilibrium during the sample period. The negative mean of the jump component could be seen as smaller downward adjustment in world crude oil demand following a large upward jump in oil price. However, the downward adjustment in demand was short-lived; the drift component of the diffusion process was very high for daily data, indicating that oil demand was pushed up by a strong income effect; consequently, oil prices were under a constant pressure to move upward. These results can be easily explained in terms of the elasticities of world demand and supply for crude oil. World demand was highly elastic with respect to world income, and highly inelastic with regard to oil prices. Crude supply has been rigid, showing little sensitivity to prices. As world real GDP expanded at 4–5 percent per year during the period under study, it caused world oil demand to expand at similar rate, creating an excess demand for oil. Given the short-term inelasticity of demand and supply with respect to prices, any small excess demand for oil would cause large variation in prices. In turn, large price increases would have small negative effect on oil demand. The negative price effect, however, would be quickly dominated by a positive income effect.

IV. CRUDE OIL PRICE AS A VARIANCE-GAMMA LEVY PROCESS

The J-D model has essentially two limitations. First, it does not capture the notion of time-varying or stochastic volatility. In particular, stochastic volatility is found to have a key role in explaining skewness and leptokurtosis in financial time-series, and in explaining the skew in market implied volatilities. In this respect, skewed distribution can arise either because of correlations between asset prices and volatility shocks, or because of nonzero average jumps. Similarly, excess kurtosis can arise either from volatile volatility or from a substantial jump component. Second, the J-D model is fit to model finite large jumps, and cannot capture

infinite small jumps which are similar to small jumps in the diffusion process. With a view to capture the notion of stochastic volatility and modeling small and frequent jumps, while simplifying computational costs, many researchers (e.g., Carr et. al (2002, 2003), Carr and Wu (2004), Cont and Tankov (2004)) have proposed the use of Levy processes for modeling asset prices. Accordingly, oil prices are modeled in this Section as a Levy process.¹⁸ More specifically, oil price returns are assumed to follow a Levy process with a variance-gamma distribution. This type of model has a simple CF and is easier to estimate.

1. Definition of the Variance-Gamma process: a variance-gamma (VG) process is defined as a Brownian motion with drift α and volatility σ , i.e. $\alpha t + \sigma B_t$, where B_t is an ordinary Brownian motion, time-changed by a gamma process. More precisely, let $G = \{G_t, t \geq 0\}$ be a gamma process with mean $a = 1/\nu > 0$ and variance $b = 1/\nu > 0$.¹⁹ Let $B = \{B_t, t \geq 0\}$ denote a Brownian motion, and let $\sigma > 0$ and $\alpha \in R$; then the VG process $X^{(VG)} = \{X_t^{(VG)}, t \geq 0\}$, with parameters $\sigma > 0, \nu > 0$ and α , can be defined as $X_t^{(VG)} = \alpha G_t + \sigma B_{G_t}$. The CF is given by

$\phi_{VG}(u; \sigma, \nu, \alpha) = E[\exp(iuX_t^{(VG)})] = (1 - iu\alpha\nu + \frac{1}{2}\sigma^2\nu u^2)^{-\frac{t}{\nu}}$. The two additional parameters in the VG distribution, which are the drift of the Brownian motion, α , and the volatility of the time change, ν , provide control over skewness and kurtosis, respectively. Namely, when $\alpha < 0$, the distribution is negatively skewed, and vice versa. Moreover, larger values of ν indicate frequent jumps and contribute to fatter tails. The moments of the log-price returns under $VG(\sigma, \nu, \alpha)$ are: the mean = α ; the variance = $\sigma^2 + \nu\alpha^2$; skewness = $\frac{\alpha\nu(3\sigma^2 + 2\nu\alpha^2)}{(\sigma^2 + \nu\alpha^2)^{3/2}}$;

¹⁸ A Levy process (LP) $(X_t)_{t \geq 0}$ has a value $X_0 = 0$ at $t = 0$ and is characterized by independent and stationary increments, and stochastic continuity, i.e., discontinuity occurs at random times. The CF of a LP is given by the Levy-Khintchine formula:

$$\phi(u) = E[e^{iuX_1}] = \exp\left(i\alpha u - \frac{\sigma^2}{2}u^2 + \int_{R \setminus \{0\}} [e^{iux} - 1 - iux 1_{|x| < 1}(x)]\nu(dx)\right), u \in R, t \geq 0. \text{ Where}$$

$\alpha \in R$ is the drift parameter, $\sigma^2 \geq 0$ is the volatility parameter, and ν is a Levy measure on $R \setminus \{0\}$, which measures jumps of different sizes. A LP is characterized by its triplet (α, σ^2, ν) .

¹⁹ The probability density of the Gamma process with mean rate t and variance νt is well known: $f(x) = x^{\frac{1}{\nu}-1} e^{-\frac{x}{\nu}} / \nu^\nu \Gamma(\frac{t}{\nu})$. Its Laplace transform is $E[\exp(-uG_t^\nu)] = (1 + \nu u)^{-\frac{t}{\nu}}$.

It results that the VG process has a simple CF $\phi_{VG}(u) = 1/(1 - i\alpha\nu u + \frac{\sigma^2\nu}{2}u^2)^{\frac{t}{\nu}}$.

and kurtosis $= 3(1 + 2\nu - \nu\sigma^4(\sigma^2 + \nu\alpha^2)^{-2})$. Clearly, skewness is influenced by α , and kurtosis by ν .

2. Estimation of the Variance-Gamma process

Let the crude oil price S_t be modeled as $S_t = S_0 \exp[\mu t + X_t]$ where X_t is a VG process. The log price return is $x_t = \Delta \log S_t = \mu + X_t - X_{t-1}$. The CF of the log price return is

$$E[\exp(iux_t)] = E[\exp(iu(\mu + X_t))] = \exp(iu\mu)E[\exp(iuX_t)] = \exp(iu\mu)(1 - i\alpha\nu u + \frac{1}{2}\sigma^2 u^2 \nu)^{-\frac{1}{\nu}}$$

Using the data described in Section II, the parameters were estimated using the ECF approach (Table 3). The estimated parameters of the VG process were stable and statistically significant, and corroborated the findings for the J-D process. Namely, when modeled as a VG process, crude oil prices exhibited a high drift coefficient, high volatility, frequent and large jumps, and skewness. The drift coefficient, estimated at $\hat{\mu} = 0.09$, asserted the presence of a strong upward pulling force which kept reigniting oil prices. The volatility parameter, $\hat{\sigma} = 1.66$, was high and significant, showing that oil markets were facing high uncertainties regarding future movements in prices. The parameter ν , which controls for tail fatness, estimated at $\hat{\nu} = 0.80$, was high and significant, implying the presence of frequent and large jumps in oil prices. The parameter α , which controls for skewness, estimated at $\hat{\alpha} = -0.36$, was negative and significant, showing that the VG distribution was left-skewed. More specifically, in response to large positive jumps in oil price, there seemed to be a cooling off period during which world crude oil demand might slowdown, causing small and frequent negative jumps in prices. However, the income demand elasticity was much higher and more significant than the price elasticity; faster world economy growth kept, therefore, pushing world oil demand upward.

Table 3. Parameter Estimates of the VG process 1/

Drift μ	Drift α	Volatility σ	Variance of VG ν
0.09	-0.146	1.66	0.804
(t=3.52)	(t= -2.54)	(t= 80.4)	(t=52.3)

1/ Using ECF method.

V. OPTION PRICING USING CHARACTERISTIC FUNCTIONS

Option Value in the Asset Price Space: Under martingale pricing, the value of an option, denoted by $f(S, t)$, is a convolution of a discounted pay-off function with the state price density. For a given final condition (pay-off): $f(S, T) = g(S)$ for all S , the option value is:

$$f(S_t, t) = E_t^Q[e^{-r(T-t)}g(S_T) | S_t = S] = \int_0^\infty e^{-r(T-t)}g(S_T)p(S_T | S_t)dS_T$$

Where Q is the risk-neutral measure. The conditional expectation is computed with respect to a risk-neutral transition probability density $p(S_T | S_t)$. However, for many stochastic processes involving stochastic volatility, jumps, or Levy type processes, transition densities are often complicated and may not be readily available in closed form. In contrast, the CF of the underlying stochastic process may be readily available in closed form. It is defined as

$$\phi(u, t) = \int_0^\infty e^{i u S_T} p(S_T | S_t) dS_T, \quad u \text{ is the transform variable.}$$

Option Value in the Fourier Space: knowledge of the CF enables to compute option prices in the Fourier space according to two alternative methods. The first method, proposed by Heston (1993), relies on a numerical inversion of the CF. However, noting the singularity of Heston's formula at $u = 0$, Carr and Madan (1999) proposed, instead, a numerical inversion of the Fourier transform of the option value. More specifically, define the Fourier transform

of the option value as $\hat{f}(u) = \int_{-\infty}^\infty e^{i u S} f(S) dS$. If $\hat{f}(u)$ can be explicitly expressed in terms of

$\phi(u)$ as $\hat{f}(u) = \hat{f}_\phi(\phi(u))$ a fast Fourier transform (FFT) inversion of $\hat{f}_\phi(\phi(u))$ would then

compute the option value from its transform as $f(S, t) = \frac{1}{2\pi} \int_{-\infty}^\infty e^{-i u S} \hat{f}(u) du$.

Let $s_t = \ln S_t$ be the log-price; $k = \ln(K)$ the log strike price; $C_T(k) =$ value of a T – maturity call option with strike K ; and $c_T(k) \equiv \exp(ak)C_T(k)$ for $a > 0$, the damped option price. The

CF of $s_t = \ln S_t$ under the risk-neutral measure is given by $\phi(u, t) = \int_{-\infty}^\infty e^{i u s_T} p(s_T | s_t) ds_T$. Let

$\psi_T(u) = \int_{-\infty}^\infty e^{i u k} c_T(k) dk$ be the Fourier transform of $c_T(k)$. Carr and Madan (1999) showed that

$\psi_T(u)$ can be expressed in terms of $\phi_T(u)$ as: $\psi_T(u) = \frac{e^{-rT} \phi_T(u - (a+1)i)}{a^2 + a - u^2 + i(2a+1)u}$. Knowledge of

the CF $\phi_T(u)$, which is the CF of the log of the asset price under the risk neutral-measure,

implies knowledge of the Fourier transform of the value of the option $\psi_T(u)$. The option

price can therefore be computed via Fourier inversion as:

$$C_T(k) = \frac{\exp(-ak)}{2\pi} \int_{-\infty}^\infty e^{-i u k} \psi(u) du = \frac{\exp(-ak)}{\pi} \int_0^\infty e^{-i u k} \psi(u) du .$$

Option pricing requires evidently knowledge of the risk-neutral process and the CF associated with this process. Under the risk neutral process, money market account discounted asset prices are martingales and it follows that the mean rate of return on the asset under this probability measure is the continuous compounded risk-free interest rate r . If the asset price S_t is modeled as $S_t = S_0 \exp[rt + X_t]$ where X_t is an LP, to obtain a risk-neutral

process verifying the martingale property, define:

$$S(t) = S(0) \frac{\exp(rt + X(t))}{E[\exp(X(t))]} = S(0) \frac{\exp(rt) \exp(X(t))}{E[\exp(X(t))]}, \text{ then } E(S(t) / \exp(rt)) = S(0). \text{ The}$$

resulting risk-neutral process for the log price is:

$$\log S(t) = (\log S(0) + rt - \log E[\exp(X(t))] + X(t)). \text{ The CF of the log price is:}$$

$$E[\exp(iu \log(S(t)))] = \exp(iu((\log S(0) + rt - \log E[\exp(X(t))]) + E[\exp(iuX(t))])). \text{ For the VG}$$

model, the resulting risk-neutral process for the asset price is:

$$S_t = S_0 \exp[rt + X_t(\sigma, \alpha, \nu) + \omega t], \quad t > 0, \text{ where, by setting } \omega = \left(\frac{1}{\nu}\right) \ln(1 - \alpha\nu - \frac{1}{2}\sigma^2\nu), \text{ Madan}$$

et al. (1998) showed that the CF for log of S_t is:

$$\phi_t(u) = \exp[\ln(S_0) + (r + \omega)t] (1 - i\alpha\nu u + \frac{1}{2}\sigma^2 u^2 \nu)^{-\frac{t}{\nu}}.$$

To obtain option prices, one can analytically invert $\phi(u)$ to get the density function and then integrate the density function against the option payoff as in Heston (1993). Alternatively, the Fourier transform of the option value can be numerically inverted using FFT as in Carr and Madan (1999). The Fourier inversion can be approximated discretely via an N -point sum with a grid spacing of Δ in the Fourier domain. The inversion integral can be approximated using an integration rule, such as Simpson's or the trapezoidal rule, as

$$\int_0^{\infty} e^{-ixu} \psi(u) du \approx \sum_{j=0}^{N-1} e^{-i\left(\frac{2\pi}{N}j\right)u_j} \tilde{\psi}_j \Delta.$$

The points u_j are equidistant with grid spacing Δ , $u_j = j\Delta$. The value of Δ should be sufficiently small to approximate the integral well enough, while the value of $N\Delta$ should be large enough to assume the CF is equal to zero for $u > \bar{u} = N\Delta$. In general, the values $\tilde{\psi}_j$ are set equal to $\tilde{\psi}_j = \psi(u_j)w_j$, where w_j are the weights of the integration rule.

VI. DENSITY FORECAST OF CRUDE OIL PRICES: THE INVERSE PROBLEM

An application of the above analysis to crude oil options is undertaken in this section with the objective of estimating, from observed options' market values, density forecast for crude oil prices at a given maturity date. The estimation of the risk-neutral distribution is known as the inverse problem in option pricing models. While the pricing problem is concerned with computing option values given model's parameters, the inverse problem consists of backing out the parameters describing risk-neutral dynamics from observed prices. The computation of a risk-neutral distribution could be seen as estimating market's expectations for future prices; in contrast, estimation of statistical distribution from realized data could be seen as the actual distribution of historical prices. Assuming a VG distribution for the log price, the inverse problem can be stated as finding the parameters $\theta = (\alpha, \sigma^2, \nu)$ by minimizing the quadratic pricing error:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{N} \sum_{j=1}^N \left(C_j^*(T, K_j) - C_j(T, K_j) \right)^2, \quad j = 1, 2, \dots, N$$

under the put-call parity constraint: $S_0 + P_j(T, K_j) - C_j^*(T, K_j) = K_j e^{-rT}$

where $C_j^*(T, K_j)$ denotes the call option computed from the VG distribution, $C_j(T, K_j)$ and $P_j(T, K_j)$ denote the observed prices of call and put options for maturity T and strikes K_j ,

respectively. $C_j^*(T, K_j)$ is given by FFT; namely $C_j^*(T, K_j) = \frac{\exp(-ak_j)}{\pi} \int_0^{\infty} e^{-iuk_j} \psi_T(u) du$.

The addition of the put-call parity condition brings extra-sample information which helps to regularize the estimation problem.²⁰ Taking into account the put-call parity constraint and choosing a penalty parameter $h > 0$, the minimization problem becomes

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{N} \sum_{j=1}^N \left(\left(C_j^*(T, K_j) - C_j(T, K_j) \right)^2 + h \left(S_0 + P_j(T, K_j) - C_j^*(T, K_j) - K_j e^{-rT} \right)^2 \right)$$

The inverse problem was applied for the VG model only for space limitation. The same methodology applies identically to the J-D model.²¹ The observed data set was for July 21, 2006; it consisted of call and put futures options contracts maturing end-September 2006; the risk-free interest rate, taken here to be the three-month U.S. Treasury bill rate, was equal to 4.965; and the crude futures price, was equal to US\$74.43/bl. The constrained minimization yielded the following triplet for the risk-neutral distribution: $\hat{\sigma}^2 = 1.72$, $\hat{\nu} = 1.12$, $\hat{\alpha} = 0.37$ which described market's expectations on July 21, 2006 regarding futures prices end-September 2006. Clearly, market participants did not anticipate any short-term change in the underlying fundamentals characterizing oil markets. They expected oil prices to remain highly volatile ($\hat{\sigma}^2 = 1.72$) and dominated by a jump process ($\hat{\nu} = 1.12$). They also expected oil prices to remain under pressure, as they assigned higher probabilities for oil prices to rise

²⁰ Cont and Tankov (2004) argued that the inverse problem could be an ill-posed problem and proposed the use of relative entropy, which is the Kullback-Leibler distance for measuring the proximity of two equivalent probability measures, as a regularization method with the prior distribution estimated from the statistical data via the maximum likelihood method. This regularization will enable to find a unique martingale measure.

²¹ The risk-neutral CF for any asset price model is given:

$E[\exp(iu \log(S(t)))] = \exp(iu((\log S(0) + rt - \log E[\exp(X(t))])E[\exp(iuX(t))])$. For the J-D model, the resulting risk-neutral process for the asset price is:

$$\phi_T(u) = \exp[\ln(S_0) + (r + \omega)T] \exp\left(T \left[-\frac{\sigma^2 u^2}{2} + i\mu u + \lambda \left(\exp\left(i\beta u - \frac{\delta^2 u^2}{2}\right) - 1 \right) \right]\right)$$

where $\omega = \left[\frac{\sigma^2}{2} + \mu + \lambda \left(\exp\left(\beta + \frac{\delta^2}{2}\right) - 1 \right) \right]$.

above the futures price level than to fall below this level. This was shown by a right-skewed risk-neutral distribution ($\hat{\alpha}=0.37$).

VII. CONCLUSIONS

Oil prices have been on a rising spree during the recent past, reaching unexpected territories, and seemed to become unbounded. Despite the importance of these prices, little was known about their underlying stochastic process. This paper studied the dynamics of oil prices during January 2, 2002–July 7, 2006. Main findings were that these dynamics were dominated by frequent jumps, causing oil markets to be constantly out-of-equilibrium. While oil prices attempted to retreat following major upward jumps, there was a strong positive drift which kept pushing these prices upward. Volatility was high, making oil prices very sensitive to small shocks and to news. The findings for both the J-D and VG specification were fully consistent with the underlying fundamentals of oil markets and world economy. More specifically, faster world economy growth during the sample period and highly expansionary monetary policies caused demand for crude oil to expand at similar pace. In view of the price inelasticities of oil demand and supply, any small excess demand (supply) would require a large price increase (decrease) to clear oil markets; hence, the observed high intensity of jumps and the strong drive for oil prices to rise.

Attention was not only limited to historical dynamics of oil prices, but it was also extended to gauging market expectations regarding future developments in these prices. Based on call and put option prices on July 21, 2006 and for maturity end-September 2006, the implied risk-neutral distribution was right-skewed, indicating that market participants maintained higher probabilities for prices to rise above the expected mean, given by the futures price, rather than fall below this mean. The risk-neutral distribution was also characterized by high volatility and high kurtosis, indicating that market participants were expecting prices to remain highly volatile and dominated by frequent jumps.

The findings of this paper could be relevant for policy-makers and industry analysts. They established the nature of the stochastic process underlying oil prices and the importance of components driving this process. An explanation of the process parameter estimates in terms of the underlying fundamentals for the oil markets was offered in order to comprehend the economics underpinning the observed oil prices dynamics. Namely, a change in the process parameters would require a change in the underlying fundamentals. Alternative modeling approaches in the paper were highly relevant for forecasting, risk management, derivatives pricing, and gauging market's sentiment; they allowed to ascertain robustness of estimated parameters. The findings of the paper could also be relevant for the Fund in monitoring oil markets and seeking policies for stabilizing these markets.

Annex 1. Method of Cumulants of Probability Distributions

Suppose that X is a real random variable whose real moment generating function is defined as $M(u) = E(e^{uX}) = \int_{-\infty}^{\infty} e^{uX} f(X) dX$, where $f(X)$ is the probability density of X . Just as the moment generating function M of X generates its moments, the logarithm of M generates a sequence of numbers called cumulants. The cumulants κ_n of the probability density of X are given by $M(u) = E(e^{uX}) = 1 + \sum_{n=1}^{\infty} \frac{m_n u^n}{n!} = \exp\left(\sum_{n=1}^{\infty} \frac{\kappa_n u^n}{n!}\right)$

Where $m_n = E(X^n)$ is the moment of order n of X . The left-hand side of this equation is the moment-generating function, so $\kappa_n/n!$ is the n th coefficient in the power series representation of the logarithm of the moment-generating function. The logarithm of the moment-generating function is therefore called the cumulant-generating function, written as:²²

$\log(M(u)) = \sum_{n=0}^{\infty} \frac{\kappa_n u^n}{n!}$. The method of cumulants attempts to recover a probability

distribution from its sequence of cumulants. In some cases no solution exists; in some other cases a unique solution, or more than one solution, exists. The relationship between moments and cumulants is of paramount importance in the estimation of the unknown parameters of the density function. First, consider moments about 0, which can be written as $m_j = E(X^j)$, $j = 0, 1, 2, \dots$. The cumulant/moment theorem says that if X is a random variable with n moments m_1, m_2, \dots, m_n , then X has n cumulants $\kappa_1, \kappa_2, \dots, \kappa_n$, and the cumulants are

related to the moments by the following recursion formula:²³ $\kappa_n = m_n - \sum_{j=1}^{n-1} \binom{n-1}{j-1} \kappa_n m_{n-j}$

Note that $m_0 = 1$. By carrying the recursion formula, the relation between raw moments and cumulants can be stated as:

$$m_1 = \kappa_1$$

²² The cumulants are also equivalently defined in terms of the characteristic function, which is the Fourier transform of the probability density function: $\phi(u) = E(e^{iuX}) = \int_{-\infty}^{\infty} e^{iuX} f(X) dX$. The cumulants κ_n are

$$\text{then defined as } \ln \phi(u) = \sum_{n=1}^{\infty} \kappa_n \frac{(iu)^n}{n!}$$

²³ This recursion formula is the Faa di Bruno's formula, equivalently written as

$$m_{r+1} = \sum_{j=0}^r \binom{r}{j} m_j \kappa_{r+1-j} \text{ for } r = 0, \dots, n-1$$

$$m_2 = \kappa_2 + m_1 \kappa_1$$

$$m_3 = \kappa_3 + 2m_1 \kappa_2 + m_2 \kappa_1$$

$$m_4 = \kappa_4 + 3m_1 \kappa_3 + 3m_2 \kappa_2 + m_3 \kappa_1$$

For central moments, defined by $\bar{m}_j = E\left(\left(X - E(X)\right)^j\right)$, the first moment \bar{m}_1 is zero; the relationship between moments and cumulants simplifies to:

$$\bar{m}_1 = \kappa_1 = 0$$

$$\bar{m}_2 = \kappa_2$$

$$\bar{m}_3 = \kappa_3$$

$$\bar{m}_4 = \kappa_4 + 3\bar{m}_2 \kappa_2$$

The first cumulant is simply the expected value; the second and third cumulants are respectively the second and third central moments (the second central moment is the variance); but the higher cumulants are neither moments nor central moments, but rather more complicated polynomial functions of the moments. The n th moment m_n is an n th-degree polynomial in the first n cumulants. Of particular interest is the fourth-order cumulant, called kurtosis, which can be expressed as $kurt(X) = E(X^4) - 3\left(E(X^2)\right)^2$.

Kurtosis can be considered as a measure of the non-Gaussianity of X . For a Gaussian random variable, kurtosis is zero; it is typically positive for distributions with heavy tails and a peak at zero, and negative for flatter densities with lighter tails.

References

- Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, Vol. 81, pp. 637–54.
- Bakshi, G., Cao, C., Chen, Z., 1997, "Empirical Performance of Alternative Pricing Models," *Journal of Finance*, 52, 2003–2049.
- Ball, C., A. and W. N. Torous, 1983, "A Simplified Jump Process for Common Stock Returns," *Journal of Financial and Quantitative Analysis*, 18, pp. 53–65.
- Ball, C., A. and W. N. Torous, 1985, "On Jumps in Common Stock Prices and their Impact on Call Option Pricing," *Journal of Finance*, Vol. 40, No. 1, pp. 155–173.
- Bates, D., S., 1966, "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options," *Review of Financial Studies*, Vol. 9, No. 1, pp. 69–107.
- Beckers, S., 1981, "A Note on Estimating the Parameters of the Diffusion-Jump Model of Stock Returns," *Journal of Financial and Quantitative Analysis*, Vol. 16, No. 1, pp. 127–140.
- Carr, P., H., Geman, D. Madan, and M. Yor, 2002, "The Fine Structure of Asset Returns: An Empirical Investigation," *Journal of Business*, Vol. 75, No. 2, pp. 305–332.
- Carr, P., H., Geman, D. Madan, and M. Yor, 2003, "Stochastic Volatility for Levy Processes," *Mathematical Finance*, Vol. 13, No. 3, pp. 345–382.
- Carr, P., and D. Madan, 1999, "Option Valuation Using Fast Fourier Transform," *Journal of Computational Finance*, Vol. 2, pp. 61–73.
- Carr, P., and L. Wu, 2004, "Time-Changed Levy Processes and Option Pricing," *Journal of Financial Economics*, 71, 113–141.
- Clark, P. K., 1973, "A Subordinated Stochastic Process with Finite Variance for Speculative Prices," *Econometrica*, Vol. 41, pp. 135–155.
- Chacko, G., and L. M. Viceira, 2003, "Spectral GMM Estimation of Continuous-Time Processes," *Journal of Econometrics*, 116, pp. 259–292.
- Cont, R. and Tankov, P., 2004, *Financial Modeling with Jump Processes*, (Chapman&Hall/CRC).
- Das, S., R. and R. K. Sundaram, 1999, "Of Smiles and Smirks: A Term Structure Perspective," *Journal of Financial and Quantitative Analysis*, Vol. 34, No. 2, pp. 211–239.

- Duffie, D., J. Pan, and K. Singleton, 2000, "Transform Analysis and Asset Pricing for Affine Jump-Diffusion Models," *Econometrica*, Vol. 69, No. 6, pp. 1343–1376.
- Fama, E.F., 1965, "The Behavior of Stock Market Prices," *Journal of Business*, Vol. 34, 420–429.
- Feuerverger, A. and R. A. Mureika, 1977, "The Empirical Characteristic Function and its Applications," *The Annals of Statistics*, Vol. 5, No. 1, pp. 88–97.
- Feuerverger, A. and P. McDunnough, 1981a, "On Some Fourier Methods for Inference," *Journal of the American Statistical Association*, Vol. 78, No 375, pp. 379–387.
- Feuerverger, A. and P. McDunnough, 1981b, "On the Efficiency of Empirical Characteristic Function Procedures," *Journal of the Royal Statistical Society, Series B*, Vol. 43, No 1, pp. 20–27.
- Feuerverger, A., 1990, "An Efficiency Result for the Empirical Characteristic Function in Stationary Time-Series Models," *The Canadian Journal of Statistics*, Vol. 18, No 2, pp. 155–161.
- Hamilton, J.D., 1983, "Oil and the Macroeconomy since World War II," *Journal of Political Economy*, 91, pp. 228–48.
- Hansen, L., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.
- Heston, S. L., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *The Review of Financial Studies*, Vol. 6, No. 2, pp. 327–343.
- International Energy Agency, *Oil Market Report*, September, 2006.
- International Monetary Fund, *World Economic Outlook*, September, 2006.
- Jorion, P., 1988, "On Jump Processes in the Foreign Exchange and Stock Markets," *The Review of Financial Studies*, Vol. 1, No. 4, pp. 427–445.
- Kendall, M. and A. Stuart, 1977, *The Advanced Theory of Statistics*, New York, MacMillan Publishing Company.
- Krichene, N., 2006, "World Crude Oil markets: Monetary Policy and the Recent Oil Shock," IMF working paper, WP/06/62.
- Madan, D., B. and E. Seneta, 1987, "Simulation of Estimates the Empirical Characteristic Function," *International Statistical Review*, 55, 2, pp. 153–161.

- Madan, D., P. Carr, and E. Chang, 1998, "The Variance Gamma Process and Option Pricing," *European Finance Review*, 2, 79–105.
- Madan, D. and F., Milne, 1991, "Option Pricing with VG Martingale Components," *Mathematical Finance*, Vol. 1, pp. 39–56.
- Mandelbrot, B., 1963, "New Methods in Statistical Economics," *Journal of Political Economy*, Vol. 61, pp. 421–440.
- Merton, R. C., 1976, "Option Pricing when Underlying Stock Returns are Discontinuous," *Journal of Financial Economics*, Vol. 3, pp. 125–144.
- Parzen, E., 1962, "On Estimation of a Probability Density Function and Mode," *Ann. Math. Statist.*, Vol. 33, pp. 1065–1076.
- Press, S., J., 1967, "A Compound Events Model for Security Prices," *Journal of Business*, 40, pp. 317–35.